

AD-A263 184

School of Economics 507

Carnegie Mellon University

Reproduced From
Best Available Copy

Argus, the Journal of Distribution, Unlimited

25-27 October 1992

DTIC
S APR 26, 1993 D
E

Mind Matters

Contributions to Cognitive and Computer Science in Honor of Allen Newell

DTIC QUALITY INSPECTED

Accession For	
NTIS	CRA&I <input checked="" type="checkbox"/>
DTIC	TAB <input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification <i>per Stu</i>	
By	
Distribution /	
Availability Codes	
Dist	Avail and / or Special
A-1	

25-27 October 1992

published by
The School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

93-08707



7001

93 4 22 048

"I want to welcome our Carnegie Mellon family and those of you that have travelled long distances to be with us today...to a celebration of the life of an extraordinary man. A scientist who achieved world-wide renown for his work. A man whose personal warmth and generosity touched each colleague and student with whom he worked. A man who was a mentor to a large number of us... and one who had a profound influence, not only on the field of computer science, but on the College of Computer Science here at Carnegie Mellon University. An influence that will be felt for many, many decades to come..."

—*Robert Mehrabian, President, Carnegie Mellon University*

Contents

5	Welcome to the Symposium: Mind Matters
	Allen Newell
7	Thoughts on the Symposium
	Tom Mitchell, Jill Fain Lehman, David Steier
13	Schedule of Events
	26-27 October 1992
19	Distinguished Speakers
	Profiles in Order of Presentation
29	Discussants
	Profiles in Order of Presentation
33	Demonstrations
	Descriptions and Locations
37	What is Soar
	The Soar Group
39	We Remember
	Herbert A. Simon, Richard Lewis, Raj Reddy, Jill Fain Lehman, Bonnie E. John, John E. Laird, Paul Rosenbloom, Stuart K. Card, David Klahr, A. Nico Habermann, Paul Allen Newell
61	Fairy Tales
	Allen Newell
65	Discovery-Integration Processes for Breaks
	In the spirit of Allen Newell - Enjoy the breaks cognitively
73	Map
	Oakland Vicinity
76	With Gratitude



Invitation to the Symposium: Mind Matters

Allen Newell

I invite you to *my* Symposium on Mind Matters. No, no, that can't be right. I invite you to *your* symposium. No, that's not quite right either, if truth be told. I invite you to *my friends'* symposium. It is my friends who thought it up; my friends who are giving the papers; my friends who are making the commentaries; my friends who are doing the work to make it happen. And it is my friends, by and large, who are the audience. (Welcome nonfriends all, if you happen to stray this way.) It is perhaps a little weird for he who is honored to be issuing the invitations. But among friends, anything is possible.

My friends also invented the title of the Symposium — *Mind Matters*. They believe I know what that means, so am the right person to tell you why that title captures the spirit of this occasion exactly. I shall try, though it is clearly an exercise in projection.

I am by imprinting a physical scientist. That's what my undergraduate days as a physics major at Stanford did to me. That, plus a father who was a Professor of Radiology at Stanford Medical School and who idolized the physical sciences. With the hubris common to physicists, I have always felt that I have known what good science is — it is theory cast in terms of mechanisms that describe how parts of the universe behave. With sometimes immense historical delay, these mechanisms always move towards being grounded in the larger mechanistic view of the universe. Theories always propose a view of how the universe is. They can never be effectively argued to be true, but only be brought before the bar of empirical evidence. All

the modern concern for contextualism, hermeneutics and the social determination of meaning has its point, but is a mere footnote to the massive evidence for this view of science. The overwhelming success within this framework of modern biology over the last half century has provided another major confirmation, if one is needed. Someday we will get another striking confirmation from cognitive science. Though it can be argued that we are well on our way, we still have an immense distance to go. Arguments are no match for the evidence that cognitive science does not control its subject matter the way physics, chemistry and now biology do.

It follows for me that a theory of mind is embodied in our theory of matter. Matter matters to mind. To put it yet another way, the great scientific question about mind is how it can occur in our physical universe as understood by the all enveloping scientific view. This is not reductionism — I'll take in any way it comes. Indeed, the modern computational view of mind, with its solutions to representation and intention, is distinctly not a simple reduction.

But it also is true of me that mind matters. What is the nature of mind is the great scientific question....

I have pursued these matters since the fifties, in concert and colleagueship with many scientific friends along the way. From my own personal viewpoint this has been an immensely cumulative trip, in which the pieces of the scientific puzzle gradually, though hardly completely, have revealed themselves. Many in psychology, never having experienced cumulative predictive science,

see it somewhat otherwise, with new paradigms and new questions moving to centerstage simply shifting the focus without cumulation. It has never seemed that way to me, though I have on occasion criticized cognitive psychology for the manner in which it fails to progress theoretically.

About ten years ago, in concert with John Laird and Paul Rosenbloom, *matters* seemed to come together.

—*A draft of an introduction to the
Mind Matters Symposium*
26 May 1992



Allen Newell

Welcome

Tom Mitchell

When a number of us began, over a year ago, to plan an event to honor Allen Newell and his many contributions, we asked each other what form the event should take. A panel of speakers to pay tribute to Allen? Absolutely not! Allen would not stand for that. His interests lay in the science, not in tribute. That's it — a scientific symposium. The most fitting way to honor a man such as Allen — to celebrate the science that he practiced with such insight and passion, and to celebrate his great influence on it through the presentations of eminent scientists who worked with him over the years.

What would be the topic? That was easy. For Allen's career formed a decades-long journey driven by a single, fundamental question: "what is the nature of mind?" From his early work with Shaw and Simon on automated reasoning, to his most recent work attempting to build a computer model of human cognition in Soar, the question that most fascinated Allen was how the mind worked. Hence this symposium on Mind Matters. In it, you will hear about the kinds of scientific questions and theories that occupied much of Allen's intellectual life. You will hear it from leading scientists, each of whom has made their own seminal contributions to this scientific puzzle, and each of whom has had the honor of working closely with Allen at some point along the way. You will hear it also from Allen's colleagues on the Soar project, who will provide commentary on each of these presentations. No doubt as you listen you will hear Allen's influence on the ideas, and Allen's enthusiasm for getting at the truth.

So welcome, on behalf of all of Allen's friends and colleagues, to Mind Matters. Thanks on their behalf to my co-organizers for their efforts, to our distinguished speakers for their presentations, to Maya Design for donating their design talent to this project, and to the National Science Foundation and Office of Naval Research for their support of this symposium. May you take advantage of the next few days to participate in the science — Allen's science, and the science of many others — and to reflect on the fact that one person can indeed have a tremendous influence on the lives and understanding of so many.

"Psychology has arrived at the possibility of unified theories of cognition—theories that gain their power by positing a single system of mechanisms that operate together to produce the full range of human cognition. I do not say they are here. But they are within reach and we should strive to attain them."

—*Unified Theories of Cognition*,
Allen Newell, 1990

Words for Thought...

1mind \mind\ *n* [ME, fr. OE gemynd; akin to OHG *gimunt* memory; both fr. a prehistoric EGmc-WGmc compound whose first constituent is represented by OE *ge-* (perfective prefix) and whose second constituent is akin to L *ment-*, *mens* mind, *monere* to remind, warn, Gk *menos* spirit, *mnasthai*, *minneskesthai* to remember — more at CO-] 1 : RECOLLECTION, MEMORY *<keep that in ~>* *<time out of ~>* 2 **a** : the element or complex of elements in an individual that feels, perceives, thinks, wills, and esp. reasons **b** : the conscious mental events and capabilities in an organism **c** : the organized conscious and unconscious adaptive mental activity of an organism 3 : INTENTION, DESIRE *<she changed her ~>* 4 : the normal or healthy condition of the mental faculties 5 : OPINION, VIEW 6 : DISPOSITION, MOOD 7 **a** : a person or group embodying mental qualities *<the public ~>* **b** : intellectual ability 8 *cap.*, Christian Science : GOD 1b 9 : a conscious substratum or factor in the universe

2mind *vt* 1 *chiefly dial* : REMIND 2 *chiefly dial* : REMEMBER 3 : to attend to closely 4 **a** : to become aware of : NOTICE **b** *chiefly dial* : INTEND, PURPOSE 5 **a** : to give heed to attentively in order to obey **b** : to follow the orders of instructions of 6 **a** : to be concerned about **b** : DISLIKE *<I don't ~ going>* 7 **a** : to be careful : SEE *<~ you finish it>* **b** : to be cautious about *<~the broken ring>* 8 : to give protective care to : TEND *~ vi* 1 : to be attentive or wary 2 : to become concerned : CARE 3 : to pay obedient heed or attention *syn* see OBEY, TEND — **mind-er** *n*

1matter \mat-er\ *n* [ME *matere*, fr. OF, fr. L *materia* matter, physical substance, fr. *mater*] 1 **a** : a subject under consideration **b** : a subject of disagreement or litigation **c** *pl* : the events or circumstances of a particular situation **d** : the elements that constitute material for treatment in thought, discourse, or writing **e** : an element of a field of knowledge, inquiry, or specialization *<~s of faith>* **f** : something to be proved in law **g** *obs* : sensible or serious material as distinguished from nonsense or drollery **h** (1) *obs* : REASON, CAUSE (2) : source esp. of feeling or emotion **i** : a condition affecting a person or thing usu. unfavorably *<what's the ~>* 2 **a** : the substance of which a physical object is composed **b** : material substance that occupies space and has weight, that constitutes the observable universe, and that together with energy forms the basis of objective phenomena **c** : a material substance of a particular kind or for a particular purpose *<vegetable ~>* **d** : material discharged from the human body 3 **a** : the indeterminate subject of reality; *esp* : the element in the universe that undergoes formation and alteration **b** : the formless substratum of all things which exists only potentially and upon which form acts to produce realities 4 : a more or less definite amount or quantity *<a ~ of 10 years>* 5 **a** : something written or printed **b** (1) : set type (2) : text material *esp*. as distinguished from illustrations 6 : MAIL 7 Christian Science : the illusion that the objects perceived by the physical senses have the reality of substance — **for that matter** : so far as that is concerned — **no matter** : without regard to : irrespective of *<was calm no matter what the provocation>*

2matter *vi* 1 : to be of importance : SIGNIFY

Words From a Friend

Jill Fain Lehman

WHEN DOES A SCIENTIFIC LIFE START?

Maxim—The scientific problem chooses you,
you don't choose it

Not at age 17—When I wanted to be a forest ranger
Not at age 19—When I wanted to be "a scientist"
Not at age 22—When I wanted to be an optical engineer
Not at age 27—When I wanted to be an organizational scientist
But at age 27.7—On a Fri afternoon in mid Nov 1954

—*Desires and Diversions*
A. Newell, December 1991

When we sat down to consider how to organize this symposium we confronted a difficult problem. Approached from an historical perspective, Allen's career has been characterized by its breadth. Focussing on the last ten years, however, his all-consuming passion (other than the Steelers) has been Soar. Granted, a Soar symposium might have been a welcome relief to the Soar community (released at last from the strictures of twelve minute presentations), but to the community at large it would have seemed at best parochial, at worst a criminal waste of an opportunity to honor breadth with breadth. How, then, to satisfy the desires of the honorers while creating more than a diversion for the honoree?

A compromise was reached (it remains to be seen if it was a workable one). Speakers would be invited to give talks in areas spanning the full historical spectrum of Allen's career. Their only responsibility would be to do good science, while the responsibility for relating the speakers' ideas to Soar would rest on discussants drawn from the Soar community. Those of us on the committee who were also members of the Soar project applauded the ingenuity of this solution. We could, no doubt, rely on Allen to be the first and best discussant for each paper, jumping in and finding any commonalities we failed to notice. Now while it is true that Allen won't be here here to help unify the seemingly disparate cognitions of our speakers and discussants, in retrospect I do not think that had we known he would be absent we would have changed our design. For the ideas will be presented, the science will continue, and links to Soar will be forged. And that is almost as good.

Information Processing and Behavioral Science

The role of the digital computer as a consequence-generating device was described....

An information processing system is specified by a set of symbols — the program. This specification contains all the implications that exist about the system, and the problem is to extract them. The computer lets us find some of these implications by being able to perform the program for any particular situation.

This is not an unfamiliar role for machines to occupy if we consider the use of computers in physical theory or statistics. However, machines have also played a different role vis-a-vis theories of man and his nature — the role of analogue. As mentioned earlier, one approach to the field of complex information processing is to construct thinking machines, and this leads naturally to the supposition that the thinking machine is to be made into a model of man. This idea is reinforced from the cybernetic movement, which in fact has used both electrical circuits with feedback and digital computers as analogies to the nervous system.

It should be evident by now that we are not using the digital computers as an analogy. It is true historically that many of the basic ideas for representing information processing systems were first formulated in connection with digital computers. It is also true that the large high speed computer is the only non-biological system currently known that can show anything like man's capability for complex behavior. But in fact we could write programs, investigate their properties, and use them to represent theories of human behavior even if computers never existed.

...I have claimed that if we have a program that behaves like a man, then we have an explanation of the behavior. I have indicated that techniques exist for specifying programs in a scientifically useful way. These share many of the virtues of the best techniques known for the description of dynamic systems — differential equations — although they are comparatively deficient in analytic power. Also, although more by indirection, I have indicated that theories of man centered around the concept of limited rationality are handled naturally by these concepts and techniques of information processing. Finally, I have shown that we are not simply drawing crude analogies between men and machines.

— *Information Processing: A New Technique for the Behavioral Sciences*, Ph.D. Thesis, A. Newell, 1957

Observations on a Scientist

David Steier

I did not know Allen Newell, except by reputation, before I came to CMU; I met him when I suppose many grad students at CMU first meet their advisors, during the information overload that constitutes the Immigration Course, when we desperately tried to figure out which advisors would be the best ones to help us change the world with our research. In my own quest, I met with Elaine Kant and Allen Newell, who described a project they had just begun, eventually known as Designer, whose stated goal was to build an AI system that would design algorithms automatically! Soon, I was hooked: here was a project that combined psychology (we studied human designers) and computer science (the design problems we studied were taken from computational geometry, a very new field then) while exploring such intriguing issues as creativity and visual imagery. I harbored dreams of building a system that would pass the algorithms portion of the theory qualifier before I did (unfortunately, this never happened, despite my best efforts to delay passing theory as long as possible).

And so I began my research, mostly working with Elaine, but with occasional guidance from Allen. I soon learned that as far as Allen Newell was concerned, there was no such thing as a disciplinary boundary beyond which his students could trust that the material "wouldn't be on the test". The Designer system grew (even including a system to aid in protocol analysis!) to the point where Elaine, Allen, and I used the services of several undergraduate programmers to keep it under control. One of these undergraduates was Brian Milnes (who later became a key member of the Soar sys-

tems staff). By 1984, he and I had built a control strategy into Designer based on a simplified form of Soar's universal subgoaling. So when Elaine Kant moved to Schlumberger and Allen became my sole advisor in 1985, we decided to use Soar itself as a basis for further work on Designer. Thus began my thesis work, leading to the creation of three Soar-based algorithm design systems. It was while I was doing my thesis research that I worked most closely with Allen, and learned firsthand of Allen's tremendous dedication to his work and his students. I remember how, during a time of particularly deep crisis, Allen and I met daily for two months for an hour or two each day, which I had assumed would be impossible for a faculty member of his stature. Even after my graduation, as I was struggling with all the dilemmas about life and work of a newly minted Ph.D., he provided constant guidance, combining a tremendous personal sensitivity with the technical depth and breadth I had grown to expect from Allen. In retrospect, I see now that I was so taken by Allen's enthusiasm for the research and the ideas that I neglected to give serious consideration to styles very different from his. But I had the chance to observe one of the great scientists in action, and I wouldn't trade the experience for the world.

"Listen to the architecture!"

—A. Newell

Workshop Report

In June 1972 an intensive workshop was held at CMU under the auspices of the Mathematical Social Science Boards (MSSB), in which a number of cognitive psychologists explored and used some new computer techniques for doing research on cognitive processes...

Since the mid-fifties there has been continuous development in experimental psychology on the view that man is a processor of information... This growth has taken a variety of forms, from providing the conceptual framework within which to ask experimental questions, to constructing computer programs to simulate specific aspects of cognitive behavior. The computer has always played a dual role in this: as a major source of the conceptual ideas about information processing systems and as a means for constructing such systems and exploring their behavior by simulation...

Two concerns provided the impetus for the present workshop. One is the continued growth, in complexity and sophistication, of the computer programs used for the study of information processing. This includes direct simulations of cognitive behavior, basic studies in artificial intelligence, and programs that aid the analysis of cognitive data. The characteristics of these current programs appear to be that: (1) each is an embodiment of some specific psychological content; (2) each permits substantial variation and modification; (3) each has a language of interaction which gives it some of the flavor of a programming language, but a language that speaks directly in psychological terms; (4) each is interactive, so that the user modifies and explores an existing system, rather than creating something from scratch; and (5) each is a large program.

The second concern is the continued ineffectiveness of scientific communication on the content of these various programs... One difficulty is that the underlying technical systems of computer science are not fully assimilated by the psychological community... But the difficulties lie also in the nature of the systems — their size, complexity and the detailed knowledge necessary for understanding and evaluation...

The goals of the workshop were to examine the existing state of the computer technology for cognitive research, working within an information processing approach. The task was not

to introduce notions of man as an information processor to those still largely unfamiliar with it, but to introduce new tools to those already working within the field. The major thesis behind the workshop — working with actual programs rather than talking about them — posed the main problems of organization and preparation.

A more important issue was the selection of programs. Only a few could be included, but this implied an undesired emphasis on the specific ones selected, for the goal of the workshop was to convey in depth the whole range of things that could be done. The selection problem was compounded by the requirement that the programs run at a given site (CMU), which implied that the assemblage of programs would exhibit undue provincialism. Any attempt to be every eclectic posed difficult problems of getting complex programs up and running on foreign computer systems...

We initially started with a set of about a dozen programs, intending to eliminate some along the way. The final number was seven...PAS-I (Protocol Analysis System I), PAS-II (Protocol Analysis System II), MAPP (Memory and Perception Program), PSG (Production System, version G), SHRDLU, CLS (Concept Learning Systems), ATS (Semi-automatic Transcription System)... The seven systems were written in six and a half languages: Stanford Lisp, Mac Lisp, Snobol 4, Fortran, SAIL (Stanford AI Language), L* (a CMU system building language) and APCOL (the language on the laboratory computer). In addition, all the programs have specialized sublanguages for run-time interaction...

...we decided to incorporate all of the programs within a submonitor-like system which would buffer the participants both from variations in language and style inherent in seven separate major programs and from the PDP10 monitor systems. This system, which came to be called ZOG, was the major programming effort designed to make communication possible... The name is not acronymic for anything. It was chosen to be short, pronounceable, capable of personification (e.g., "ZOG told me that ..."), and non-conflicting with other names.

— *Report on a Workshop in New Techniques in Cognitive Research, Carnegie Mellon University, 21-29 June 1972,*
A. Newell, H.A. Simon, R. Hayes, L. Gregg

Schedule of Events

Monday, 26 October 1992

Carnegie Mellon Research Institute
4400 Fifth Avenue
Main Auditorium - 2nd Floor

8:00 a.m.	12:00 p.m.
Continental Breakfast <i>Auditorium Foyer</i>	LUNCHEON, Ballroom The Holiday Inn at University Center 100 Lytton Avenue
8:45 a.m.	1:30 p.m.
Welcoming Remarks: Robert Mehrabian President, Carnegie Mellon University	Problems in Sense Resolution — <i>George A. Miller</i>
8:50 a.m.	2:15 p.m.
The Evolution of the Soar Cognitive Architecture — <i>John E. Laird</i>	Discussant: Jill Fain Lehman
9:45 a.m.	2:30 p.m.
The Human Cognitive Architecture: Why and How It Can Be Studied — <i>Zenon N. Pylyshyn</i>	Constraints on Processing Capacity: Architectural or Implementation? — <i>Marcel A. Just and Patricia Carpenter</i>
10:30 a.m.	3:15 p.m.
Discussant: Richard Lewis	Discussant: Richard Young
10:45 a.m. BREAK	3:30 p.m.
11:00 a.m.	Personal Reflections Members of the audience to participate.
Integrating Inductive and Analytic Learning — <i>Tom M. Mitchell</i>	4:30 p.m.
11:45 a.m.	Soar Project Demonstrations Presentation of "Desires and Diversions" Software Engineering Institute 4500 Fifth Avenue, adjacent to CMRI
Discussant: Paul Rosenbloom	

Monday Evening, 26 October 1992

6:30 p.m.

COMMENCEMENT OF SHUTTLE SERVICE
TO THE CARNEGIE SCIENCE CENTER
Continuous service from each hotel
6:30 pm - 7:15 pm

7:00 p.m.

RECEPTION

THE CARNEGIE SCIENCE CENTER
One Allegheny Avenue, North Side
Parking Available in Center Lot

8:30 p.m.

Introduction

**The Impact of Allen Newell on AI
and Computer Science**
— *Raj Reddy*

ON VISITING THE CARNEGIE SCIENCE CENTER...

Situated on the Ohio River near the Pittsburgh Point and Three Rivers Stadium, the science center was created as a symbol of the city's commitment to science education in the coming decade. The center is the fifth component of The Carnegie, which also includes The Carnegie Library system, Carnegie Music Hall, Scaife Art Museum, and The Carnegie Museum of Natural History in Oakland (near Carnegie Mellon University). The Science Center also incorporates the former Buhl Science Center, the 1939 science building located on the North Side of Pittsburgh in what is now called the Allegheny Square Annex. This latter site continues to serve as an education facility for public courses, special programs and community outreach activities.

Space Places, an extraordinary record of astronomical exploration, features the work of noted space photographer Roger Ressmeyer. He has spent more than 20 years traveling to space observation centers around the world while on assignment for magazines such as National Geographic, Life, Time, Newsweek, and Smithsonian. From Kublai Khan's 13th century observatory in Beijing to the giant Keck Telescope atop Hawaii's Mauna Kea, Ressmeyer has compiled a portrait of how far we have traveled and what lies ahead in space exploration. The exhibit showcases the human, mystical and technical sides of space exploration, with emphasis on internationally strategic space places such as optical observatories, space probe communications facilities, launch sites, satellite preparation centers, and astronaut and cosmonaut training academies. Taken on location in the U.S., French Guiana, Chile, Australia, the former USSR, Egypt, China, Japan and Europe, the photographs highlight space technologies and historic centers of the international space age.

For Another Visit

The Carnegie Science Center, which opened on October 5, 1991, is four floors of exhibits and educational programs covering almost every field of science — lasers, computers and robotics, nutrition, water movement, industrial science, the ocean and much more. The Center has three theaters: the Rangos Omnimax Theater features larger-than-life films on a colossal 79-foot-diameter domed screen; the Henry Buhl, Jr. Planetarium and Observatory enables visitors to chart their own star paths in the largest physically-interactive planetarium in the world; the Blue Cross Health Science Theater uses multimedia presentations to examine the latest in health medicine. The Works, an Electronic Exhibit Theater, is an interactive, demonstration and presentation area all rolled into one. It combines industrial shows, complete with lasers, molten metal and live electric arcs. The environment combines open exhibits visitors can explore themselves, regularly scheduled demonstrations, and a 20-minute extravaganza, the Science Spectacular, which includes a working foundry, robots, cryogenics, electrical areas and lasers. The Center also houses a self-sustaining aquarium, a science "pier" for students, and a mechanized model of the digestive system. Today is only a welcome. Come back again.

Aquabatics shoots two stories of good, clean fun straight through an opening in the ceiling. Designed especially for the Science Center by WET Design of Universal City, California, this interactive water sculpture enables visitors to control water patterns and jets using computerized stations. The exhibit illustrates important hydraulic principles such as laminar flow and Bernoulli's Principle as the water travels up and down the sculpture. If visitors are not operating the fountain, computers take over the direction of the patterns.

Science Way seems like any other city street, but visitors soon encounter many surprises on this byway that *disguises* scientific principles within common objects and events. In this small neighborhood composed of two houses, a park and a mechanic's shop, scientists have changed everyday items into fascinating science experiments that encourage visitors to discover the processes scientists use to understand phenomena experienced in their daily lives.

USS Requin, a World War II diesel-electric submarine, has been explored by over 150,000 people since her arrival in Pittsburgh in fall 1990. A blend of science and history, the sub tour highlights the various types of sonar equipment installed aboard Requin, enabling visitors to distinguish the different types of sonar and to discover how that technology allowed submarines to navigate and operate underwater. The Requin has an array of ingenious voice, visual and electronic communications equipment on board. Tour guides, many of whom are former submariners, demonstrate voice-activated telephones, decipher the international signal flag alphabet, and explain electronic detection systems.

Venus: Earth's Fiery Twin is on view in the Henry Buhl, Jr. Planetarium & Observatory and features images of Venus from the Magellan space probe. *The Dream is Alive* is playing in the Rangos Omnimax Theater with in-flight footage shot by 14 astronauts during three space shuttle missions in 1984.

— Courtesy of The Carnegie Science Center

Problems of Communication

A language is Janus-like: one face toward each communicant. When both are human and substantially identical one is hardly aware of the two-faced nature. When there is great inequality — a mother to her child, a boy to his dog, a man to a computer — we view the language only in terms of the weaker. The problem is to get the computer to understand the instructions we give it; the man can shift for himself. We have already partially learned our lesson with respect to computers; the whole development of programming languages reflects an attempt to ease the task of the human. But at each step the course is roughly the same: a new linguistic feature is proposed (for human benefit) and the problem is how to construct the systems to interpret it and use it efficiently. All the science is built around the machine side of the language problem; the "new linguistic features" come of themselves, full blown, out of the minds of those experienced in the programming art, or those struggling with a problem.

Still, the human is an information processing system in his own right with his own special properties and limitations. At some point, when our ability to construct programming languages becomes sufficiently great, we will need to know a great deal more about the human and his communication problems. What are his languages? How does he interpret linguistic expressions? What stages of gradual refinement and definition do his intentions go through on the way to becoming a set of instructions to be given to a computer? What is the role of mnemonic symbols for him, and how complex can they become before they cease to be helpful? All these questions are fundamentally psychological and linguistic, and they have interest in their own terms. Our concern with them here is as applied psychology — as the human engineering of language and communication with the computer.

Let us give two examples of work in this area that we have an interest in pursuing, and where we have given some preliminary consideration to the problems. The first concerns on-line continuous communications (so called conversational mode) between man and computer. We are currently pushing this problem, and like everyone else we are proceeding by designing systems, putting them into operation, discovering the difficulties and revamping the system. But it is clear that the human has difficulties thinking clearly in real-time — witness the difference between spoken and written English. Ultimately (after the first exploitation of the new powers opened up by the hardware), the communication language must adapt to these limitations of the human. One could believe this would happen eventually just by successive approximation. One could proceed a good deal faster by understanding in some depth the nature of the human limitations, and then with this understanding in hand, designing language systems to reflect it.

The second example stems from searching for alternatives to on-line interaction: Perhaps the human could simply communicate his rough plans to the machine, and let the computer fill in the details. This puts the computer in the role of problem solver again, but the problems would not be very difficult (that is what it means to be a "detail"). However, before one can worry about the computing program, one needs to know a good deal about the nature of human plans — about what it is that the human can communicate while he is still vague about the details.

— *Proposal for a Center for the Study of Information Processing*, A. Newell, A. Perlis, E. Schatz, 1964

Tuesday, 27 October 1992

**Carnegie Mellon Research Institute
Main Auditorium - 2nd Floor**

8:15 a.m.	12:15 p.m.
Continental Breakfast <i>Auditorium Foyer</i>	LUNCHEON, Auditorium Greek Community Center 143 North Dithridge Street
9:00 a.m.	1:45 p.m.
Big Ideas in Computer Structures <i>—Gordon Bell</i>	Children, Adults and Machines as Discovery Systems <i>— David Klahr</i>
9:45 a.m.	2:30 p.m.
Discussant: Milind Tambe	Discussant: Tony Simon
10:00 a.m.	2:45 p.m.
Software Architectures for Integrating Independent Systems <i>— Mary M. Shaw</i>	The Mind as a Reasoning Machine <i>— Philip N. Johnson-Laird</i>
10:45 a.m.	3:30 p.m.
Discussant: David Steier	Discussant: Thad Polk
11:00 a.m. BREAK	3:45 p.m. BREAK
11:15 a.m.	4:00 p.m.
The Human, the Machine, the Task and Their Interaction <i>— Stuart K. Card</i>	The Patterned Matter that is Mind <i>— Herbert A. Simon</i>
12:00 a.m.	5:00 p.m.
Discussant: Bonnie John	CONCLUSION
	Thank you for joining us.



Distinguished Speakers

In Order of Presentation

ROBERT MEHRABIAN

Robert Mehrabian is the seventh President of Carnegie Mellon University. He assumed this position July 1, 1990. Prior to his appointment at Carnegie Mellon he served from 1983 to 1990 as Dean of the College of Engineering, University of California at Santa Barbara.

Dr. Mehrabian, is an internationally recognized materials scientist. He is credited with leading UCSB's engineering college to national prominence. During his tenure, the college initiated a doctoral program in materials science, recruited more than 65 new faculty and established seven research centers in areas such as robotics, high-performance composites, compound semiconductors, and risk studies and safety in chemical and nuclear plants.

Mehrabian joined UCSB after four years in the federal government, first as chief of the Metallurgy Division, later as director of the Center for Materials Science at the National Bureau of Standards (now NIST).

At the NBS, he initiated major industry/government programs, including phase diagram programs with the American Society for Metals, the American Ceramics Society and the Society for Plastic Engineers, and data programs with the National Association of Corrosion Engineering and the Welding Research Council. In addition, Mehrabian started a program with the American Iron and Steel Institute to develop advanced sensors for the steel industry.

From 1975-1979, Mehrabian was a professor of metallurgy and professor of mechanical engineering at the University of Illinois at Urbana-Champaign. He began his academic career at the Massachusetts Institute of Technology (MIT) in 1968 as a research associate and later was promoted to assistant professor.

For the past seven years, he has been a member of

ALCOA's Science and Technology Advisory Council. Mehrabian was one of the founders of Superconductor Technologies Incorporated, a Santa Barbara based company which produces devices using new high-temperature superconductors. In addition, Mehrabian has been a consultant to numerous companies in the field of materials, and has had a long-term relationship with Pratt & Whitney Aircraft, developing new materials for advanced aircraft engines.

He has published more than 130 technical papers and edited six books in metallurgy and materials science. He holds eight U.S. and over thirty foreign patents as co-inventor of new metalworking and composite fabrication processes.

In 1991, The Metallurgical Society awarded the Leadership Award to Mehrabian. In addition, he is a fellow of the American Society for Metals International and an elected member of the National Academy of Engineering. In 1983, he received the Henry Marion Howe Medal of the American Society of Metals. In 1980, he was the George Kimball Burgess Memorial Lecturer.

Mehrabian is of Armenian descent. Born in Tehran, Iran, he is an American citizen and a 1960 graduate of Phillips Exeter Academy in Exeter, New Hampshire. He received the B.S. (1964) and Sc.D. (1968) in Metallurgy at the Massachusetts Institute of Technology.

JOHN E. LAIRD

John Laird is Associate Professor of Electrical Engineering and Computer Science at the University of Michigan. He received his B.S. from the University of Michigan in 1975 and his Ph.D. in Computer Science from Carnegie Mellon University in 1983. Before joining the faculty at the University of Michigan, Dr. Laird was a member of

the research staff at Xerox Palo Alto Research Center from 1984 to 1986.

Dr. Laird's research career began in 1976 when he started graduate school at Carnegie Mellon University, where his primary research interests were in the nature of the architecture underlying artificial and natural intelligence. During the next sixteen years, he had the great pleasure of pursuing his research while working with Allen Newell, who was first his advisor, later his colleague, and always a mentor. While at CMU, Dr. Laird initially worked on the Instructable Production System project and following its demise explored the use of problem spaces as an organizing framework for production systems with Dr. Newell. These explorations led to the first Soar architecture, a general problem solving system that was able to use a wide variety of weak methods. Concurrently, Paul Rosenbloom had been working with Dr. Newell on learning mechanisms, and in the summer of 1983, Laird, Newell, and Rosenbloom formed the Soar Research Project to study general cognitive architectures.

Since 1983, research in Soar has thrived, with Dr. Laird focusing on research on learning, problem solving, and the continued development and evolution of the Soar architecture. Over those years, Laird, Rosenbloom and Newell maintained a close working relationship and jointly led the Soar Research Project.

Today and into the foreseeable future, Dr. Laird's central research focus is Soar. He is currently using Soar for cognitive modeling and the construction of intelligent agents that interact with dynamic external environments. At the center of the work on intelligent agents is the integration of the capabilities required of autonomous agents, including reactive execution, planning, learning, and natural language. His initial work involved the construction of intelligent robotic systems, and has progressed to attempting to create intelligent agents that behave with the same flexibility, adaptability, and robustness as humans in large-scale, realistic, simulated environments populated by multiple agents.

ZENON N. PYLYSHYN

Zenon Pylyshyn received a B.Eng. in Engineering-Physics from McGill University in 1959, an M.Sc. in Control Systems from the University of Saskatchewan in 1961 and a Ph.D. in Experimental Psychology from the University of Saskatchewan for research involving the application of information theory to studies of human short-term memory. Following his Ph.D. he spent two years as a Canada Council Senior fellow and since then has been on faculty at The University of Western Ontario in London, where he is at present Professor of Psychology and of Computer Science as well as honorary professor in the departments of Philosophy and Electrical Engineering. He is also Director of the UWO Centre for Cognitive Science. In the past several years he has also been on the faculty at Rutgers University in New Brunswick, New Jersey where he spends part of each year as director of the Rutgers Center for Cognitive Science. He is also president of CogniCom, a company consisting of seven academic researchers that provides consulting services in artificial intelligence and new generation computer technology. He is a former holder of a Killam Fellowship and is currently a Fellow of the Canadian Institute for Advanced Research (CIAR).

Dr. Pylyshyn has been a consultant on data processing to the government of Saskatchewan (Psychiatric Services Branch), a consultant on artificial intelligence to the U.S. National Institute of Health, a consultant to Xerox Palo Alto Research Center, a consultant to the Alfred P. Sloan Foundation and to the Computer Science and Engineering Research study of the U.S. National Science Foundation. He is a regular consultant to various granting agencies, including the Social Science and Humanities Research Council, the Natural Science and Engineering Research Council, the National Institute of Health, the National Science Foundation and the National Endowment for the Humanities. In addition he is a regular editorial consultant for about a dozen journals, a member of the editorial boards of the journals Cognitive Science, Cognitive Psychology, Cognition, Computational Intelligence, Medical Expert Systems, Mind and Language, and Artificial Intelligence, and associate editor of Behavioral and Brain Sciences, and Cognitive Neuroscience. He is a founding member and a member of the governing boards of the Canadian Society for Computational Studies of Intelligence and the International Cognitive Science Society, a Fellow of the Canadian Psychological Association, a member of the advisory board of the Cognitive Neuroscience Institute of New York, past president of the Society for

Philosophy and Psychology, and past president of the Cognitive Science Society. He also serves on several technical advisory boards, including the Department of Communications, National Research Council of Canada Committee on Artificial Intelligence, B.C. Advanced Systems Institute, and PRECARN Associates (Precompetitive Applied Research Network).

Dr. Pylyshyn is recipient of numerous fellowships and awards and has held invited visiting positions at various universities. He received the Donald O. Hebb Award of the Canadian Psychological Association in June 1990 for his contribution to psychology as a science.

Dr. Pylyshyn has delivered over 100 invited talks and seminars at various universities and research laboratories in Canada, U.S.A., Europe and Japan. He has published over 60 scientific articles and book chapters and is author of *Computation and Cognition: Toward a Foundation for Cognitive Science* (1984). He also edited and contributed to four books, including: (*Perspectives on the Computer Revolution*) (1988); *Computational Processes in Human Vision: An Interdisciplinary Perspective*) (1988). He published (*The Robot's Dilemma: The Frame Problem in Artificial Intelligence*) (1987) and (*Meaning and Cognitive Structure: Issues in the Computational Theory of Mind*) (1986). As chairman of a panel on artificial intelligence under a grant from the National Science Foundation Dr. Pylyshyn also helped to produce a major survey of the state-of-the-art in artificial intelligence which appeared in the book *What Can be Automated* (1980).

Dr. Pylyshyn's research interests have covered a wide range of topics over the past twenty years. He began with research on performance of human operators in control systems, which extended to include experimental studies of short term memory and the time course of information transformations with rapid forgetting. Later he included projects on the computer analysis of text (particularly the content analysis of psychiatric interviews) and on the role of grammar and of intonation on the perception of sentences and of spectrally distorted speech. This work involved some theoretical studies of psycholinguistics. Other research carried out in collaboration with members of the computer science department was concerned with machine vision, and the development of computer models of human perceptual-motor coordination. He is principal investigator on several major grants and contracts.

For the past three years, Dr. Pylyshyn's personal research has dealt with two general areas. One is the theoretical analysis of the nature of the human computa-

tional system which enables humans to reason and perceive the world. This has led to a number of theoretical investigations of what is referred to as the "architecture of the mind" or the human "cognitive architecture". On the more experimental side Pylyshyn has been concerned with exploring his so-called FINST theory dealing with human visual attention and the encoding of spatial information. This theory consists of a set of hypotheses about a mechanism by which the location of features in a visual display are preattentively indexed (i.e., "FINSTed") so that they can be referred to by subsequent cognitive processes. Several papers have been published on this theory and its experimental investigation, as well as its implications for understanding perceptual-motor coordination and for the design of human interfaces.

TOM M. MITCHELL

Tom M. Mitchell is a Professor of Computer Science and Robotics at Carnegie Mellon University, and is Director of the Design Systems Laboratory within the Engineering Design Research Center. He earned his B.S. degree (1973) from MIT and his M.S. (1975) and Ph.D. (1978) degrees from Stanford University. He taught in the Computer Science Department at Rutgers University from 1978 until moving to Carnegie Mellon in 1986. In 1983 he received the IJCAI Computers and Thought award in recognition of his research in machine learning, and in 1984 a National Science Foundation Presidential Young Investigator Award. In 1990 he was elected a Fellow of the American Association for Artificial Intelligence. His current research focuses on artificial intelligence, machine learning, robotics, and engineering design.

GEORGE A. MILLER

George Armitage Miller was born February 3, 1920, in Charleston, West Virginia. His parents were divorced in 1927 and he and his mother returned to Charleston where they lived with her parents while he attended public schools. He graduated from Charleston High School in 1937, then followed his mother and her second husband to Washington, D. C., where he lived with them and attended the George Washington University for one year.

When his step-father was transferred to a regional Office in Birmingham in 1938, Miller transferred to the University of Alabama. As an undergraduate he studied history and, partly to overcome adolescent shyness, became active in Blackfriars, the student drama club, where he met Katherine James—who became his wife in 1939. Interests in acting and stage production were supplemented by courses in the Speech Department, and led to a double undergraduate major in History and Speech. In 1940 he received a Bachelor of Arts degree.

The following year he was a graduate student and teaching assistant in the Speech Department at Alabama, where he took courses in phonetics, voice science, and speech pathology—courses that persuaded him that he needed to learn psychology in order to contribute to the field. In 1941 he received a Master of Arts in Speech and was awarded a fellowship to continue in Speech at the University of Iowa, but had to refuse because he was unable to afford the cost of moving his family. He was rescued from this dilemma by D. A. Ramsdell, the Professor of Psychology at Alabama, who offered him a position teaching psychology. Miller served as an Instructor in Psychology from 1941 to 1943.

In 1942 Miller attended Summer School at Harvard, where he returned as a full-time graduate student in 1943. Although he had intended to study clinical psychology at Harvard, his knowledge of speech and hearing had prepared him for research in psychoacoustics. Consequently, during World War II he worked on military voice communications as Research Associate under S. S. Stevens at the Harvard Psycho-Acoustic Laboratory. At war's end, Miller submitted some of his military research on speech perception as a dissertation, and received his Ph.D. in Psychology from Harvard University in 1946.

At Harvard after the war, first as a Research Fellow in the Psycho-Acoustic Laboratory and then as an Assistant Professor in the Department of Psychology, Miller continued and extended his studies of speech production and perception. In 1948 C. E. Shannon's mathematical theory of communication provided conceptual tools that Miller needed in order to understand data that had accumulated during and following the war: in particular, the inverse relation between the intelligibility of a speech signal and the number of alternative signals that might have occurred instead. Shannon's measure of the amount of selective information in a message inspired a series of experiments measuring how far a listener's expectations influenced his perceptions. Miller summarized that work in 1951 in *Language and Communication*, a text that helped to establish psycholin-

guistics as an independent field of research in psychology.

From 1951-55, as Associate Professor of Psychology at the Massachusetts Institute of Technology, Miller continued this research on speech perception, and tried to extend Shannon's measure to account for short-term memory as well. The observation that the uncertainty of a message has far less effect on a person's ability to remember it than on his ability to perceive it led Miller to propose that short-term memory is limited, not by the number of bits of information it can hold, but by the number of "chunks" into which the message is recoded. In 1956 the implications of the hypothesis that a person can retain about seven chunks in short-term memory were developed in a widely quoted (and often misquoted) paper, "The Magical Number Seven, Plus or Minus Two."

In February 1955 Miller returned to Harvard's Department of Psychology. His attempt to estimate the amount of information per word in conversational speech led him to the linguist, Noam Chomsky, who showed him how the sequential predictability of speech follows from adherence to grammatical, not probabilistic, rules. The next decade was spent developing and testing the psychological implications of Chomsky's grammatical theories. Some of those ideas found expression in 1960 in *Plans and the Structure of Behavior*, a book written jointly with E. Galanter and K. Pribram. In 1960 Miller was co-founder, along with J. S. Bruner, of the Harvard Center for Cognitive Studies, and in 1964 he was appointed Chairman of the Department of Psychology.

Miller visited The Rockefeller University in New York in 1967, and in 1968 decided to stay there as Professor of Experimental Psychology. By then his research interests had shifted from syntax to lexicon, and in 1976 *Language and Perception*, written with P. N. Johnson-Laird (in large measure during a visit to the Institute for Advanced Study in 1971-72) presented a detailed hypothesis about the way lexical information is stored in a person's long-term memory. Back at Rockefeller, Miller attempted to test some aspects of the hypothesis with studies of the development of language in young children; that project was summarized in 1977 in *Spontaneous Apprentices: Children and Language*.

In 1979 Miller moved to Princeton University, where he is now James S. McDonnell Distinguished University Professor of Psychology, Emeritus. In 1986, in collaboration with Gilbert Harman, he established the Princeton Cognitive Science Laboratory. In 1990 he wrote *The Science of Words*.

He resides with his wife, Katherine, in Princeton, New Jersey. They have two grown children: Nancy and Donald James, and three grandsons.

MARCEL A. JUST PATRICIA CARPENTER

A few months after Marcel Just received his undergraduate degree in psychology from Montreal's McGill University in 1968, and Patricia Carpenter received hers from the University of Iowa, they met while enrolling in the graduate program in the psychology department of Stanford University. Over the next four years at Stanford, they learned many valuable research skills in experimental and cognitive psychology.

After joining the faculty at Carnegie Mellon University in 1972, they used these skills to study language comprehension and reading. Their work has evolved around several related themes. They generally try to provide a process account of what occurs psychologically while a person is reading or solving a problem. The empirical methodologies they use (such as eye-tracking) are particularly appropriate for process-tracing. The research has often attempted to account for the variation among individuals, in terms of the processes that distinguish one person from another. The resulting theories often take the form of computational models. Their work has provided an account of the psychological processes in reading, in solving intelligence test items, and in visual thinking.

RAJ REDDY

Dr. Raj Reddy is the Herbert A. Simon Professor of Computer Science and Robotics and Dean of the School of Computer Science at Carnegie Mellon University. From 1979 to 1992, also served as the Director of the Robotics Institute at Carnegie Mellon University. Prior to joining Carnegie Mellon's Department of Computer Science in 1969, Dr. Reddy was an Assistant Professor of Computer Science at Stanford University from 1966 to 1969. He also served as an Applied Science Representative for International Business Machines Corporation (IBM) in Australia from 1960 to 1963.

His research interests include the study of human-computer interaction and artificial intelligence. His cur-

rent research projects include speech recognition and understanding systems; multimedia presentation technology; collaborative writing, design and planning; the World Language Bank project; and the Automated Machine Shop project.

His professional honors include: Fellow of the Institute of Electrical and Electronics Engineers; Fellow of the Acoustical Society of America; Fellow of the American Association for Artificial Intelligence; Member of the National Academy of Engineering; and President of the American Association for Artificial Intelligence, 1987-89. Dr. Reddy was presented the Legion of Honor by President Mitterrand of France in 1984.

GORDON BELL

Gordon Bell is a computer industry consultant at large. He spent 23 years at Digital Equipment Corporation as Vice President of Research and Development, where he was responsible for Digital's products. He was the architect of various mini- and time-sharing computers and led the development of DEC's VAX and VAX Computing Environment. Subsequently he was a founder and Vice Chairman of Encore Computer Company, and a founder of Stardent Computer. Bell has been involved in the design of about 50 computers, a dozen of which were multiprocessors, and many products at Digital, Encore, Stardent, and a score of other companies. Bell is a director of The Bell-Mason Group that provides systems for venture development.

During 1966-72 he was Professor of Computer Science and Electrical Engineering at Carnegie Mellon University. He worked with Allen Newell on research that led to the classic book *Computer Structures: Readings and Examples* (1971). *Computer Structures* posited the Processor-Memory-Switch, and Instruction-Set Processor notations for describing computers. *Computer Structures: Readings and Principles* (1982) by Siewiorek, Bell, and Newell further refined the notations. The taxonomy for Computer Structures has proved to be valid throughout the various computer generations, and equally descriptive of the various approaches for massively parallel computers of the 1990s.

Allen's interest in design helped establish the framework for *Designing With Register Transfer Modules* (1972) by Bell, Grayson, and Newell. This work at the Register Transfer Level allowed the first fully automated computer design systems to be pioneered at CMU. In the

early 1990s, commercial design systems are finally beginning to synthesize computer systems.

In 1986-1987 he was the first Assistant Director of the National Science Foundation's Computing Directorate. He led the National Research Network panel, and was an author of the High Performance Computer and Communications Initiative.

Mr. Bell has co-authored six books and many papers about computer structures and most recently, start-up companies. In April 1991, Addison-Wesley published *High Tech Ventures: The Guide to Entrepreneurial Success*, which describes the Bell-Mason Diagnostic, for analyzing new ventures. Gordon is a member of various professional organizations including the AAAS (Fellow), ACM, IEEE (Fellow and Computer Pioneer), and the National Academy of Engineering. He is on the boards of Cirrus Logic, Chronologic Simulation, University Video Communications, and Visix Software, and is a technical advisor to various companies including Kendall Square Research and Microsoft. He is a founder and director of The Computer Museum, Boston.

His awards include: The Mellon Institute Medal, and IEEE's McDowell and Eckert-Mauchly Awards, and the Von Neumann Medal. President Bush awarded Bell The National Medal of Technology in 1991 "for his continuing intellectual and industrial achievements in the field of computer design; and for his leading role in establishing cost-effective, powerful computers which serve as a significant tool for engineering, science, and industry."

Mr. Bell is married to Gwen Bell, the founding president of The Computer Museum. He resides in Los Altos, California and Boston, Massachusetts.

MARY M. SHAW

Mary Shaw is Associate Dean for Professional Education Programs and a Professor of Computer Science at Carnegie Mellon University. She has been a member of the faculty since completing the Ph.D. degree at Carnegie Mellon in 1972. From 1984 to 1987 she served as Chief Scientist of CMU's Software Engineering Institute, where she still holds a joint appointment. She had previously received a B.A (cum laude) from Rice University and worked in systems programming and research at the Research Analysis Corporation and Rice University.

Her research interests in computer science lie primarily in the areas of programming systems and software engineering, particularly software architecture, programming languages, specifications, and abstraction techniques. Particular areas of interest and projects have included software architectures (Vitruvius), technology transition (SEI), program organization for quality human interfaces (Descartes), programming language design (Alphard, Tartan), abstraction techniques for advanced programming methodologies (abstract data types, generic definitions), reliable software development (strong typing and modularity), evaluation techniques for software (performance specification, compiler contraction, software metrics), and analysis of algorithms (polynomial derivative evaluation).

Dr. Shaw is an author or editor of six books and more than seventy papers and technical reports. She is a Fellow of the IEEE and American Association for the Advancement of Science, and a member of the Association for Computing Machinery, the Society of the Sigma Xi, and the New York Academy of Sciences. She serves on the National Research Council's Computer Science and Telecommunications Board, Working Group 2.4 (System Implementation Languages) of the International Federation of Information Processing Societies, and the IEEE Technical Committee on Software Engineering. In addition, she has served on a number of advisory and review panels, conference program committees, and editorial boards.

STUART K. CARD

Stuart Card is a Research Fellow and Manager of User Interface Research at the Xerox Palo Alto Research Center. He received his A.B. in Physics from Oberlin College in 1966 and came to Carnegie Mellon University to study with Allen Newell and Herb Simon under the Systems and Communication Sciences program. At CMU he pursued an interdisciplinary program in psychology and computer science with Allen Newell as his thesis advisor, resulting in a Ph.D. in psychology in 1978. He has been an adjunct faculty member at Stanford University and an affiliate at CMU.

With Newell as consultant, Card and Thomas Moran (another Newell student) founded the Applied Information Processing Psychology project at Xerox PARC to work on developing psychological models that could be used as a basis for engineering practical systems. One result was the book with Newell and Moran, *The*

Psychology of Human-Computer Interaction. Card's research interests have continued to center around the theory and design of interactive computing systems. His studies of input devices led to the Fitts's Law characterization of the mouse and was one of the factors leading to its commercial introduction. His characterization of window system use in terms of spatial working sets led to the development (with Austin Henderson) of the Rooms multiple workspace window manager. He is co-editor of the book, *Human Performance Models for Computer-Aided Engineering*.

Currently, his research is concerned with the design and analysis of interactive 3D user interfaces for information access systems and on interfaces using computational paper. With George Robertson and Jock Mackinlay, he has developed the Information Visualizer system, a system for information visualization and retrieval. He is also a co-author of the recent *ACM SIGCHI Curriculum for Human-Computer Interaction*.

DAVID KLAHR

David Klahr received his Ph.D. in 1968 from Carnegie Mellon's Graduate School of Industrial Administration in Organizations and Social Behavior. His dissertation advisor was Herbert Simon. From 1967-69, he was an Assistant Professor at the University of Chicago with a joint appointment in the School of Business and the Department of Mathematics. He returned to CMU with a joint appointment in GSIA and Psychology in 1969, and became Professor of Psychology at CMU in 1975, and Head of the Department of Psychology in 1983.

Throughout his career Dr. Klahr's research has focused on the analysis of complex cognitive processes in such diverse areas as voting behavior, college admissions, consumer choice, peer review, problem solving and scientific reasoning. He pioneered the application of information processing analysis to questions of cognitive development, and, in collaboration with Iain Wallace, formulated the first production system models to account for children's performance on a variety of tasks. His current research interests include the acquisition of quantitative skills, and the development of scientific reasoning skills.

Dr. Klahr was a visiting research Fellow at the University of Stirling (Scotland) in 1968 and a Fulbright Lecturer at the London Graduate School of Business Studies in 1969. He is a Fellow of the American Psychological

Association, and a member of the Society for Research in Child Development, the Cognitive Science Society, and the Psychonomic Society. He has served on the Editorial Board of several cognitive science journals and on the subcommittee on Memory and Cognitive Processes for the National Science Foundation. In addition, to being Head of Carnegie Mellon University's Department of Psychology, he is also Director of its Center for Scientific Literacy.

PHILIP N. JOHNSON-LAIRD

Philip N. Johnson-Laird is a Professor in the Department of Psychology at Princeton University. He received his B.A. (Hons, First Class) in Psychology from the University College London in 1964 and his doctorate in Ph.D. in Psychology in 1967. He earned an M.A. from Cambridge University in 1988 and was awarded an Honorary Doctorate from the University of Goteborg in 1983.

Dr. Johnson-Laird was Assistant Director of the Applied Psychology Unit, the Medical Research Council, Cambridge from 1982-1989. Prior to this appointment he was a member of the Experimental Psychology group at the University of Sussex, as a reader (1973-1978) and then as professor (1978-1982). Among other accomplishments, Dr. Johnson-Laird served as a freelance musician and music critic for talks on the BBC Third Programme from 1960-1961.

Dr. Johnson-Laird was the recipient of the Rosa Morison Memorial Medal from University College London in 1964, where he also earned the James Sully Scholarship. He has been awarded the Spearman Medal of the British Psychological Society (1974), the Presidents' Award of the British Psychological Society (1985), and the Medaglia d'Onore of the University of Florence (1989). He was made a Fellow of Darwin College, Cambridge (1985-1989) and continues to serve as a Fellow of the British Academy.

He is the recipient of numerous grants, including: the Medical Research Council grant for research into the effects of linguistic variables on cognitive performance, in conjunction with Dr. P. C. Wason, 1968-1971; Social Science Research Council grants for scientific assistance for research into the representation of meaning in the mental lexicon, 1973-1976; into the role of inference in comprehension and memory, in conjunction with K. Ehrlich, 1977-1979; for an experimental investigation into linguistic performance, 1977-1978; for an investiga-

tion into an integrated model of syntactic, semantic and inferential processing, 1978-1981; for an investigation for working memory and inference with Jane Oakhill, 1982-1983; and a Science Research Council Computing Grant with H.C. Longuet-Higgins and N.S. Sutherland, 1980-1984.

Dr. Johnson-Laird has been a member of many learned societies, including the British Psychological Society, the Linguistics Association, the Experimental Psychology Society, a founding member of the Cognitive Science Society, the Association for Computational Linguistics, the International Pragmatics Association, the International Society for Research on Emotion, and the Society of Experimental Psychologists. He has also served on numerous research councils: the Psychology Committee, Social Science Research Council (PB) 1975-1979; the Linguistics Committee, Social Science Research Council (PB) 1980-1982; the Advisory Council of the International Association for the Study of Attention and Performance, 1984; and the Education and Social Research Council, Committee to set up Interdisciplinary Research, Centre on Language, 1988 (Chair).

He has served on the editorial boards of the Quarterly Journal of Experimental Psychology (1970-1972), Cognitive Psychology (1975-), Cognition (1975), Behavioral & Brain Sciences (1976-), Cognitive Science (1977-87), Cognition & Brain Theory (1980-5), Journal of Semantics (1982-), Cognitive Development Abstracts (1983-5), Language and Cognitive Processes (1985-), Journal of Experimental Psychology: General and Computer Speech and Language (1985-), and the American Journal of Psychology (1989-). He is the author of over 130 publications in psychological, linguistic and cognitive science journals and has authored six books, including: *Psychology of Reasoning* (with P.C. Wason), *Language and Perception* (with George A. Miller), *Mental Models*, and *The Computer and the Mind: An Introduction to Cognitive Science*.

Since 1955, he has been engaged in research on artificial intelligence and computer simulation of human thinking. He is a recipient of the Nobel Prize in Economics, the Distinguished Scientific Contributions Award of the American Psychological Association, the Turing Award (with Allen Newell) of the Association for Computing Machinery, the Gold Medal for Research of the American Psychological Foundation, and the National Medal of Science, and has been awarded honorary degrees by a number of American and foreign universities.

For forty years, Simon has been a frequent and close collaborator with Allen Newell, beginning with studies of information processing in air defense operations, and continuing with their early explorations (from 1955) in artificial intelligence which led to the Logic Theorist, GPS, and their book, *Human Problem Solving*. Severally and separately, they have advanced the claim of physical symbol systems to be the natural habitats of human and computer intelligence.

HERBERT A. SIMON

Herbert A. Simon is Richard King Mellon University Professor of Computer Science and Psychology at Carnegie Mellon University, having joined the University in 1949 with the founding of the Graduate School of Industrial Administration. He was educated (B.A. and Ph.D.) at the University of Chicago, and had done research and taught at the University of California (Berkeley) and Illinois Institute of Technology before coming to Carnegie.

Carnegie Mellon University Professor Receives National Medal of Science

FOR RELEASE: June 23, 1992

PITTSBURGH—Allen Newell, the U.A. and Helen Whitaker professor of computer science at Carnegie Mellon University, has been awarded a National Medal of Science for his pioneering work in artificial intelligence, the theory of human cognition, and development of computer software and hardware systems for complex information processing.

He will be honored by President Bush on June 23, along with seven other distinguished scientists, in a ceremony and dinner at the White House.

The National Medal of Science, authorized by Congress in 1959, is given periodically by the President in special recognition of outstanding contributions to knowledge in the physical, biological, mathematical or engineering sciences. Some 304 medals have been awarded since President Kennedy named the first recipient in 1962.

Newell's career spans the entire computer era, which began in the early '50s. Newell is known as one of the founders of both artificial intelligence and cognitive science. Both fields grew in part from his idea that computers could process symbols as well as numbers, and if programmed properly, would be capable of solving problems in the same way that humans do.

In cognitive science, he has focused on problem solving and the cognitive architecture that supports intelligent action in humans and machines. In computer science, he has worked on areas as diverse as list processing, computer description languages, hypertext systems and psychologically based models of human-computer interaction.

Since the early 1980s, Newell's work has centered on the development of Soar, an artificially intelligent software system capable of solving problems and learning in ways similar to human beings. The goal of the Soar project is to provide an underlying structure that would enable a computer system to perform a range of cognitive tasks. The system has been in use for the past five years as the framework for several intelligent systems at research institutions around the country.

A native of San Francisco, Newell received a bachelor's degree in physics from Stanford University in 1949. He spent a year at Princeton University doing graduate work in mathematics, and worked for the Rand Corporation as a research scientist from 1950-61. While at Rand, he met Nobel Laureate Herbert A. Simon, then a professor of industrial administration at Car-

negie Institute of Technology (CIT), now Carnegie Mellon University. Their discussions on how human thinking could be modeled led Newell to earn a doctor's degree in industrial administration from CIT's business school in 1957.

Dr. Newell joined the Carnegie Mellon faculty as a professor in 1961. He has played a pivotal role in elevating its School of Computer Science to world-class status.

Newell has written and co-authored more than 250 publications, including 10 books. He co-authored "Human Problem Solving" with Herbert Simon in 1972, and co-authored "The Psychology of Human-Computer Interaction" with two other colleagues in 1983. His most recent book, "Unified Theories of Cognition," published by Harvard University Press in 1990, is based on the thesis that tools are at hand that will allow cognitive scientists to develop one unified theory to describe many different types of behavior, instead of building separate theories to cover isolated aspects, as has long been the practice. A system based on a unified theory could support the full range of intelligent behaviors.

Newell's awards and honors include the Harry Goode Award of the American Federation of Information Processing Societies (1971); the A.M. Turing Award of the Association for Computing Machinery, joint with Simon, (1975); the Alexander C. Williams Jr. Award of the Human Factors Society (1979); the Distinguished Research Contribution Award of the American Psychological Association (1985); the Research Excellence Award of the International Joint Conference on Artificial Intelligence (1989); the Emanuel R. Piore Award of the Institute for Electrical and Electronic Engineers (1990); and the Franklin Institute's Louis E. Levy Medal (1992). He was awarded honorary doctoral degrees by the University of Pennsylvania and Groningen University (The Netherlands).

Newell is a member of the National Academy of Sciences, the National Academy of Engineering and the American Academy of Arts and Sciences. He was the first president of the American Association for Artificial Intelligence. In 1987 he delivered the Williams James Lectures to the Department of Psychology at Harvard. These lectures formed the basis for his book, "Unified Theories of Cognition."

Newell is married and lives in Pittsburgh with his wife Noel. The couple's son, Paul Allen, lives in California.



Discussants

BONNIE JOHN

"I was Allen's last non-Soar student": this was the defining fact of Bonnie's graduate school life.

She arrived at that distinction through a circuitous route. Starting intellectual life as a mechanical engineer, Bonnie received a Bachelors of Engineering from The Cooper Union in NYC (1977) and a Masters of Science in Mechanical Engineering from Stanford University (1978). She worked as a mechanical designer, then as a systems engineer writing specifications for data- and tele-communications systems at Bell Laboratories.

In her role as a specifications writer, Bonnie wanted the human factors psychologists at Bell Labs to tell her if the system she was designing was going to be easy for people to learn and use. They responded, "You design it, we'll test it." In other words, she had to give them a prototype to test before they could tell her anything about the system. Naturally, this was often too late in the product cycle to make substantial changes to the design. She countered, "You don't build a bridge and then see if it falls down." So she went looking for psychological theories to help her put human performance considerations into the early stages of system design. She looked in the Human Factors program at Stevens Institute of Technology (Hoboken, NJ) - but found only thousands of experiments and no usable theories. Finally, Bonnie quit my high-paying engineering job to become a slave to science (otherwise known as a CMU Psychology graduate student) and attempt to beat psychology into a shape useful for engineering design.

Newell was going on sabbatical. Newell wasn't accepting new grad students that term. Newell only accepted students to work on Soar. But, living in Psychology, Bonnie didn't hear those rumors until the hour before her first appointment with him, so she went anyway.

They were both so pleased by the coincidence of their views - his coming from academic observation of computer design, hers coming from the frustrations of doing it - that they just kept meeting, and eventually discovered that he was her advisor.

As an assistant professor in CMU's Computer Science Department, Bonnie tells us of her choice to work with Soar:

"Soar is my diversion. Allen had continued his diversion into HCI for a few more years because of me; I figured I owed him a couple of years of attention to Soar. I say Soar is on the path to engineering-style models of computer users - and it may be - but it certainly wasn't the most direct path after my thesis. It will probably be a long diversion (as all diversions are long); it's already been useful in a couple of ways; the salvage remains to be seen."

JILL FAIN LEHMAN

Jill Fain Lehman is a Research Computer Scientist in Carnegie Mellon University's School of Computer Science and the current principal investigator of the Soar project at CMU. She received a B.S. in computer science from Yale in 1981 followed by her M.S. and Ph.D. in computer science from Carnegie Mellon in 1987 and 1989, respectively.

Although her research interests have always been in natural language processing, Jill first worked with Allen Newell (and Peter Hibbard) on a what-you-see-is-what-you-get interface for interactive page layout. It was a profoundly ill-fated piece of research that convinced her never again to stray too far from her natural language pursuits. Despite the ultimate failure of their first en-

deavor together, Allen nevertheless agreed to be on Jill's thesis committee, bemoaning only that the work on a self-modifying natural language interface was not being done in Soar. Upon completing her thesis work, Jill joined the Soar project as a post-doc, working with Allen (and great trepidation) on modelling education in interactive microworlds, and with Allen and Rick Lewis on NL-Soar, the theory of natural language comprehension in Soar. She continues to work on NL-Soar, with the grand scheme of unifying the current comprehension theory with theories of generation and acquisition, and on microworld education, with the grand scheme of unifying that work with NL-Soar.

RICHARD LEWIS

Richard L. Lewis is a Ph.D. candidate in the School of Computer Science at Carnegie Mellon University. He attended the University of Central Florida in his hometown of Orlando, graduating with a B.Sc. in computer science in 1987. Through the encouragement of his professors at UCF, he developed a taste for basic research that ultimately led him to CMU.

Allen Newell first began making an impression on Rick at the 1987 CMU computer science departmental reception. Soon after, Rick scheduled a meeting in hopes that Allen might decide to take him as a student. Armed with little more than some background in graphics and a computer science degree from a less-than-well-known school (and figuring he had little to lose), Rick announced that his research interest was "finding out how the mind works", and concocted a "plan" whereby he might work on Soar while supposedly drawing upon his experience in graphics. Allen, with his characteristic daring and boldness, took Rick on as a student, promptly tossed aside his original plan, and convinced him that they should begin to tackle natural language comprehension in Soar---this despite the fact that Rick wasn't even sure what psycholinguistics was. At that point Allen undertook to educate Rick about science and cognition, a process Allen remained committed to for the next five years.

Rick's current research interests, not surprisingly, are in building computational theories of human cognitive processes, especially language comprehension. He is a member of the Cognitive Science Society, the American Psychological Society, the American Association for Artificial Intelligence, and the Association for Computational Linguistics.

THAD POLK

Thad Polk recently finished an interdisciplinary Ph.D. degree in computer science and psychology under Allen Newell and is currently starting a post-doc in neuropsychology with Martha Farah at the University of Pennsylvania. It may seem strange that one of Allen's former students is a burgeoning neuropsychologist but, in fact, it was through Allen's inspiration and encouragement that Thad decided on that path.

Dr. Polk graduated from the University of Virginia with an undergraduate degree in pure mathematics in May of 1986. He applied to do graduate work in computer science at a variety of schools with the hope of pursuing a career in artificial intelligence. A few weeks later he picked up the phone to find myself talking with Allen Newell! Despite the fact that he was a world-known scientist while he was just an *ignorant* undergraduate, he spoke to Thad as if he was already a colleague — with the utmost respect and a genuine interest in his plans and goals. Then and there he decided to go to CMU and was virtually certain he wanted to study with Allen. Those were the "two best professional decisions" he ever made.

Working with Allen inspired Thad's interest in human cognition and he consequently decided to augment his computer science education by studying psychology, leading to the interdisciplinary Ph.D. It was also through their interactions that he became particularly interested in cognitive architecture, and it is that interest that has led him to cognitive neuropsychology — a field that offers the hope of studying the architecture more directly than has been possible using traditional methods.

As should be clear, Allen Newell has been the single most important influence on his scientific development. The fact that Thad's research path is very different from those of others who have interacted with him should not be surprising. It is a tribute to the breadth of his ideas as well as their power.

PAUL ROSENBLUM

Paul S. Rosenbloom received his B.S. degree in Mathematical Sciences from Stanford University in 1976, but was left with an unsatisfied craving to understand the mind and the production of intelligent behavior. He was then drawn to Carnegie Mellon University by the charms of its computer science department and the possibility of working with Allen Newell, whom he first

met in 1976 while visiting Carnegie Mellon as a prospective graduate student. It was abundantly clear on first meeting that Dr. Newell had the vision, enthusiasm, ideas, and expertise to finally help satisfy Dr. Rosenbloom's craving. Under Dr. Newell's tutelage, Dr. Rosenbloom first worked on the Instructable Production System project, and then on models of human practice and stimulus-response compatibility. In the process, he received his M.S. and Ph.D. degrees in Computer Science from Carnegie Mellon in 1978 and 1983, respectively.

Since receiving his Ph.D., Dr. Rosenbloom continued to work closely with Dr. Newell, but now as a collaborator on the Soar project — a multi-disciplinary attempt at understanding the architecture underlying both natural and artificial intelligence — first at Carnegie Mellon (1983-1984), where he was a Research Computer Scientist; then at Stanford University (1984-1987), where he was an Assistant Professor of Computer Science and Psychology; and then at the University of Southern California (1987 -), where he is presently an Associate Professor of Computer Science and a Project Leader at the Information Sciences Institute. Within this overall research focus on architectures for intelligence, and the more specific focus on Soar, Dr. Rosenbloom has active research interests in machine learning, problem solving and planning, models of memory (and their implementation), autonomous agents in simulation environments, expert systems, neural networks, and cognitive modeling.

Dr. Rosenbloom has been Chair of the Special Interest Group on Artificial Intelligence of the Association for Computing Machinery (ACM SIGART), an editor of the journal "Machine Learning", and program co-chair of the Tenth National Conference on Artificial Intelligence (AAAI-92), and is now a Councillor of the American Association for Artificial Intelligence (AAAI).

TONY SIMON

Tony Simon is Assistant Professor in the Cognitive Science group of the School of Psychology at Georgia Institute of Technology. His research is focussed on cognitive transition mechanisms in general and the development of quantification abilities in particular. In both areas, Soar and Allen Newell's ideas about "Basic Quantitative Codes" are central components.

Born in London, England in 1960 he received his B.A. in Psychology from Lancashire Polytechnic in 1981 and Ph.D. from Sheffield University in 1984. Trained as a developmental psychologist, his primary interest was in the mechanisms of cognitive change and early in his research career, he became interested in information processing approaches to the problem. During a year each on the faculty of the London School of Economics and Lancashire Polytechnic he became convinced that computational modeling was the only way to directly study cognitive change.

In 1986 he joined Richard Young at the Applied Psychology Unit in Cambridge to work on user-models in Human Computer Interaction. On this project he was introduced to Soar. Sent to Carnegie Mellon for a month to study Soar's learning mechanism he chose as a project "transition mechanisms in development" which he had been excited to see Newell address in video tapes of the William James lectures. In 1988 Tony Simon joined the Soar project and worked with Allen Newell and David Klahr in creating Q-Soar, a theory of the development of number conservation in young children. He is currently editing a book (with Graeme Halford) on computational models of cognitive development.

DAVID STEIER

David Steier is a research scientist with the Engineering Design Research Center and the School of Computer Science at Carnegie Mellon University. He received his B.Sc. from Purdue University, West Lafayette, Indiana in 1982, his M.S. from Carnegie Mellon in 1985, and his Ph.D. from Carnegie Mellon in 1989, all in computer science. His doctoral thesis was entitled "Automating Algorithm Design within a General Architecture for Intelligence." He is co-author (with P. Anderson) of the monograph *Algorithm Synthesis: A Comparative Study*. He has also held summer research positions at Schlumberger-Doll Research, Texas Instruments, and Evanston Hospital. His current research interests center on the application of AI techniques, especially integrated problem-solving and learning architectures, to problems in engineering design.

RICHARD YOUNG

Richard M. Young studied Engineering and Artificial Intelligence (AI) before earning his Ph.D. in Psychology from Carnegie Mellon University in 1973. From 1973 to 1978 he was a Research Fellow in the Department of Artificial Intelligence at Edinburgh University. Since 1978 he has been on the scientific staff of the United Kingdom Medical Research Council's Applied Psychology Unit in Cambridge. He also acts as a consultant to Rank Xerox Cambridge EuroPARC, and is a former Chairman of the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB). His research interests lie in human-computer interaction and the use of AI for the computer simulation of human thinking, problem solving and cognitive skills. For the last few years he has been working with Soar, mainly to serve as the basis for models of the computer user. He first met Allen Newell when he worked as a research programmer for him in 1966-1967. (In those days, it was possible to learn what there was to learn about AI in a single semester!) Allen served as his Ph.D. supervisor from 1968-1973.

Each new machine that is built is an experiment. Actually constructing the machine poses a question to nature; and we listen for the answer by observing the machine in operation and analyzing it by all analytical and measurement means available. Each new program that is built is an experiment. It poses a question to nature, and its behavior offers clues to an answer. Neither machines nor programs are black boxes...we can open them up and look inside...and draw many lessons from a single experiment.

— *Computer Science as Empirical Inquiry: Symbols and Search*, ACM Turing Award Lecture 1975, A. Newell, H.A. Simon¹

¹ACM Turing Award Lectures: The First Twenty Years 1966-1985, ACM Press, 1987.

Demonstrations

Monday, 26 October 1992

On Monday, 26 October, demonstrations of select projects within Soar will be run at the Software Engineering Institute. The demos commence at 4:15 p.m. and conclude at 6:00 p.m.

Demonstrations are scheduled in two locations: (1) Conference Room 4000, 4th floor; (2) Training Room B, 2nd floor. There are eight scheduled demos, each of which will run continuously. The following pages highlight each demonstration.

The videotape *Desires and Diversions* will be presented in the Auditorium on the first floor.

**Refreshments will be served in the Cafe on the first floor
of the Software Engineering Institute from 4:15 to 6:00 pm.**

DESires AND DIVERSIONS

Videotape

Time: 4:30 pm (One viewing)

Allen Newell

Auditorium

Allen Newell delivered a School of Computer Science Distinguished Lecture on December 4, 1992. He shared his thoughts on a life of science at this, the last formal presentation he made to his colleagues, students and friends at Carnegie Mellon University. Copies of his slides are available.

Abstract: What happens with a whole scientific career? Mine isn't over, but it's already forty years long. What shape and even purpose can be given to such an endeavor? Like all things human, the answers are diverse — there are many styles of scientific lives. The one I know best, of course, is my own. It hardly seems unique to me, but it does typify a definite style. I will tell some of the story of my own total scientific endeavor, and try to tell it to shed light on how to live a science.

RedSoar

Conference Room 4000
Nasir K. Amra <amra@med.ohio-state.edu>

RedSoar is a knowledge-based abductive system that attempts to model how the domain expert solves antibody identification tasks in the blood bank domain. It embeds the generic task methodology and visual routines within Soar.

InstructoSoar

Training Room B
Scott Huffman <huffman@engin.umich.edu>

Learning from natural language instructions: this is a demo of InstructoSoar, a Soar system which learns new operators from imperative instructions given in natural language (using NL-Soar). The demo will illustrate InstructoSoar's two-stage learning process, in which initially a contextualized, "episodic" memory of the instruction content is learned, and future recall and execution of the episode results in general procedural learning.

Air-Soar

Conference Room 4000
Douglas J. Pearson <douglas@engin.umich.edu>

The demo will consist of showing Air-Soar controlling a plane as it takes off and flies a set course, and/or a demo of it landing. From the Soar point of view it's an exercise in continuous control at many levels of abstraction in real time. The system is reactive, at each level responding to unexpected events and dealing with the unpredictable nature of the domain.

Soar/Mathematica

Conference Room 4000
Dhiraj Pathak <dkp@centro.soar.cs.cmu.edu>

The demo consists of a run of Soar/Mathematica on a engineering design task that exhibits the performance and acquisition processes making up the system.

DSI

Training Room B
Thomas McGinnis <tgm@centro.soar.cs.cmu.edu>

The Developmental Soar Interface is an integrated set of menu-driven tools for writing, debugging, and explaining Soar models. It includes an interactive graphic display of models on the problem space level, and structured editors for Soar productions and TAQL programs. Its components have been installed at eight sites and are used by more than 30 scientists.

Soar/PA

Training Room B
Frank Ritter <ritter@centro.soar.cs.cmu.edu>

Discussion

Soar/PA is an extension to the Developmental Soar Interface for testing process models routinely with both verbal and non-verbal protocol data. It helps the analyst automatically or semi-automatically align the model's predictions with the protocol actions, and provides a graphical summary of where the model could be improved. The DSI can then be used to modify the model, and the cycle of testing and improving can then be repeated.

Sched-Soar

Training Room B
Josef Nerb <nerb@vax1.rz.uni-regensburg.dbp.de>

Sched-Soar tries to solve a simple job-shop-scheduling task in a cognitive plausible way. Confronted with the problem, Sched-Soar starts an analysis of the situation to find out whether there was a similar situation for which there is more information available. When insufficient knowledge can be accumulated to make a decision, Sched-Soar asks the environment (a super-intelligent agent) for help. Afterwards, Sched-Soar tries to explain why the "hint" from outside is on the path to solution. That explanation will be memorized explicitly. The explanation is hypothesis-driven and often incomplete, leading to a series of interesting transfer effects (both positive and negative).

Soar Typing Model

Conference Room 4000
Garrett Pelton <gap@centro.soar.cs.cmu.edu>

This is a Soar model of typing that is built from a hierarchical description of the task. The learned system comes close to modelling human behavior. The limits on how much it learns seems to be mainly dependent on the Soar architecture.

Desires and Diversions
SCS Distinguished Lecture
Carnegie Mellon University
4 Dec 91
Allen Newell
Computer Science
Carnegie Mellon University

Current Version: 09 December 91 10:55

TABLE OF CONTENTS

1. Title page
Preface — You may think I have to do this
2. Large Soar systems—Growth of chunks
3. Large Soar systems—Growth of the Rete net
4. Large Soar systems—No average growth effect
5. Styles of scientific lives
6. My style—The pursuit of a single desire
7. When does a scientific life start?
8. Diversions—And their uses #1
9. Diversions—And their uses #2
10. Failures—And their uses #1
11. Failures—And their uses #2
12. Failures—And their uses #3
13. Successes—And abandoned opportunities #1
14. Successes—And abandoned opportunities #2
15. Successes—And their cultivation
16. A deep idea transformed—The problem space
17. Deep ideas transformed—Brief treatment of the others
18. Choosing the final project
19. Maxims for a dedicated scientific life

— Prepared by A. Newell, 1991

What is Soar

The goal of the Soar project is the development of a uniform computational architecture that supports a wide variety of intelligent behavior and can function as the basis of an integrated intelligent agent. Scientifically, Soar plays a variety of roles. It is a vehicle for conducting research on artificial intelligence, integrating problem-solving, learning, and interaction with the environment within a system capable of autonomous operation. It is the basis for a candidate unified theory of human cognition that models a range of human cognitive functioning from immediate-response tasks, memory, learning, skill, problem solving, reasoning and language. We also intend for Soar to point the way towards the next generation of expert systems, which will be capable of both acquiring and applying substantial amounts of knowledge to perform difficult intellectual tasks.

The Soar architecture is based on formulating all goal-oriented behavior as search in problem spaces. A problem space consists of a set of states and a set of operators that move between states. A goal is formulated as the task of reaching one of a desired set of states from a specified initial state. Under conditions of perfect knowledge, satisfying a goal involves starting at the initial state, and applying a sequence of operators that result in a desired state being generated. Knowledge is represented as productions, which fire to add information to a temporary working memory when the conditions of the productions match the current problem-solving situation. When knowledge is not perfect, the system may not know how to proceed. For example, it may not know which of a set of operators should be applied to the current state. When such an impasse occurs, Soar automatically generates a subgoal to resolve the impasse. These subgoals are themselves processed in additional problem spaces, possibly leading to further impasses. The overall structure is one of a hierarchy of goals, with an associated hierarchy of

The Soar Group

problem spaces. Interaction with the external environment is carried out at the top-level space: perceptual input (e.g. from a vision system) and motor commands (e.g., to control a robot arm) are represented as augmentations of the top-level state. When a goal is terminated, the problem-solving that occurred within the goal is summarized in new productions called chunks. If a situation similar to the one that created the chunk ever occurs again, the chunk fires to prevent any impasse, leading to more efficient problem solving.

Soar was initially developed at Carnegie Mellon University by John Laird (now at the University of Michigan), Allen Newell, and Paul Rosenbloom (now at the University of Southern California). The project has grown to include over a hundred researchers distributed across multiple institutions, disciplines, and continents. Soar 5, which runs in Common Lisp, is currently available and documented for public use; we are also in the process of releasing Soar 6, a significantly faster C-based version.

Soar's behavior has been studied over a range of tasks and methods. The tasks include standard AI toy problems, routine tasks, and knowledge-intensive tasks (design, classification and planning problems). A number of Soar models have now been built that explain and predict human behavior in domains such as immediate reasoning, lexical acquisition, strategy change and cognitive development. Current areas of research include the application of Soar's problem solving and learning capabilities to control of robots and of intelligent agents in simulation environments to improve adaptability, the generation and comprehension of natural language, the integration of multiple methods for planning, the integration of knowledge learned from a variety of sources, and the use of Soar to model human-computer interaction. We are also developing efficient hardware and software technology for implementing Soar, taking advantage of the potential speedups for parallel processing and new match algorithms.



Herbert A. Simon

Last week, three days after I had written the remarks I'm about to give, I heard for the first time Allen's Distinguished Lecture of last December, which I think many of you here did hear. I heard the videotape, heard him say it all, describing his path in life, with his invariable clarity, drama and humor. From time to time there comes a man or woman who has a dream of vision, who not only dreams it, but gives body to it. Brings it to life. Allen Newell was such a man. He had a vision of what human thinking is. He spent his life enlarging that vision, shaping it, materializing it in a sequence of computer programs that exhibited the very intelligence they explained. If you asked Allen Newell what he was, he'd say, "I'm a scientist." He played that role almost every waking hour of every day for nearly half a century. As a great scientist, he was also a great artist — struggling to create form against the severe constraints imposed by nature.

Science is not science fiction. It accepts the tests of observation and experiment, acknowledges the supremacy of fact over wish or hope. The smallest experiment can crash to earth the most attractive theory. This is the art Allen Newell practiced. Modeling mind, testing his models by experiment and observation, revising them to suit the obdurate facts.

We humans have long been obsessed by four great questions: the nature of matter, the origins of the universe, the nature of life, the workings of mind. Allen Newell chose for his life work answering the fourth of these questions, explaining the human mind. That choice had already been made when I

met him in Santa Monica early in 1952 and conversed with him as he sat perched on a desk in Rand's Systems Research Laboratory. In the first ten minutes of our conversation, I knew his urge to understand the mechanisms of human thought. The great issues that occupied Allen's mind were never held secret long.

Soar as it stands just now, and *Unified Theories of Cognition*, his recent book, represent the answer Al had reached when his work was brought to an end. Not a final answer, as he would be the first to say, but an arrangement of important masses of stone in a cathedral, salient for defining the shape it will take. Allen was serious, but not solemn. Whimsy and laughter came easily and often to him. Life, sometimes perplexing, was not a plodding march but a vivid drama in which he acted with brilliance and eclat, quite aware of the dramatic effects he was producing. This too was obvious early on.

The Systems Research Laboratory, which Allen built with three colleagues, operated a human experiment on the grandest scale, simulating an air force early warning station, its task an entire air force unit. Only Allen and his co-directors could have dreamed up theory on this mega-bucks scale, at a time when behavioral scientists might timidly request five or ten thousand dollars for their research.

Understanding that great science could only be done inside nurturing institutions, Al devoted much of his energy to building and improving the environments in which he worked. First, the Systems Research Laboratory; then the burgeoning

computer simulation in the basement of the GSIA¹ building, then the Computer Science and Psychology Departments at Carnegie. In later years, the computer networking of our entire campus.

His contributions to national institutions were no less important. For example, his advisory roles in DARPA and other Washington agencies and presidencies of two young professional organizations, the American Association for Artificial Intelligence and the Cognitive Science Society.

How is Allen seen and felt as a person? Sometimes his energy, his intensity, his single-mindedness could overwhelm people, but only until they came to understand his goals. And anyone who, as Allen did, can deal with me for forty years, on a nearly daily basis, without a single quarrel — occasional scientific disagreements, of course, but no quarrels — any such person has remarkable qualities of endurance and forgiveness. Our friendship is so interwoven with our work together, that there can be no thought of the science without thought of the friendship. What a bonanza, to have such a companion to share the risk and the excitement, the effort and the pleasure of striving toward discovery. One with whom you advance with bravado, side by side, on a skeptical world. A companion as competitive and willful as yourself, but between you a synchrony of goals and heuristics that let you pursue a common path for forty years.

Allen's loyalty was immense, as was his willingness to shoulder the tasks that make organizations work. In every organization he served he acquired, as if by a law of nature, a major leadership role. Collaborating with all in an atmosphere of utmost trust and cordiality. He devoted enormous time and energy to counseling his colleagues, taking their concerns professional and personal as his, from dean and department head, to junior faculty member and newest graduate student. Allen was trusted, respected and loved because his motives were totally clear. There was no guile in him. Like all of us, he was pleased to be recognized. Receiving the National Medal of Science in his last

months gave him, I think, great satisfaction. But recognition, save the ultimate recognition of his contributions to understanding mind, was not his lodestar. He worked to enable good science to be done. He worked to advance the resources, the effectiveness, the human warmth of the organizations in which he lived.

There is much more to say about Allen as a friend, as a husband to Noël, as a father to Paul. He was all of these and these partnerships provided crucial support for his life as a scientist. He devoted himself to them at least as thoughtfully and lovingly as he did to his professional activities. But saying more about that part of Allen will have to wait a few weeks or months, until the pain of losing him has dulled a little.

A discussion of chess playing could follow this line of analysis — could concentrate on the current status of chess programs and their implications for whether machines do or do not pose a threat to the evolution of our society. But the study of chess programs can serve other interests; for example, interests in human psychology. How do humans in fact play chess or, more generally, how do they think and reason?

—*An Example of Human Chess Play in the Light of Chess Playing Programs*, A. Newell and H.A. Simon,

1964

¹Graduate School of Industrial Administration

Richard Lewis

I am Rick Lewis and I was privileged to be a graduate student of Allen's. As many of you know, during the past eight or nine years, Allen's scientific passion centered on Soar. Those of us who worked on the Soar project at CMU, as I did as his graduate student, were fortunate enough to be able to spend a great deal of time with Allen. To convey what he meant to each of us would be impossible, knowing we can only give you a glimpse.

We carry many happy memories of Allen with us. His dedication and boundless energy were both amazing and inspirational. Many of us can remember working late on the machine, sending electronic mail messages to Allen at 1:00 a.m., and receiving a response fifteen minutes later. But while we were still sleeping, Allen was back on the machine at 7:30 the next morning, working on the next proposal to keep us funded. Deadlines seemed to have little effect on his style. Ten page messages would often suddenly appear, clarifying some difficult issue when we knew he was supposed to be working on a book chapter that was due last month. Even more contagious was his enthusiasm and the great pleasure he took in the science and the people who lived it. It didn't take us long to learn the secret to a great meeting with Allen. Simply bring lots of raw data or Soar traces. He would squint his eyes, dive in with both feet, mutter *that's interesting*, and not put it aside for days to come.

He simply loved the Soar Workshops, when people from all over the world would get together and present their research on Soar. I can still see him standing at the side of the room, responding to one

of the talks with some animated discourse about problem spaces or the knowledge level, when he would be invariably cut off by the buzz of the timer, indicating that the question period was over. He would stop cold in mid-sentence and a broad grin would spread across his face. For as much as he would have enjoyed continuing his thoughts, he took even greater delight in the whole rapid fire of one scientific nugget after another, the style we had adopted for the workshops.

These facets of Allen would have been more than enough to make indelible impressions on us all. But what remains most dear to us is the way Allen related to us personally. In spite of his own scientific greatness, we were all his colleagues, from faculty to staff to graduate students. He took us seriously from day one. It mattered not whether we came from Harvard or Yale or Central Florida. He treated us all as if we had something to contribute. The community was a set of unique individuals in his eyes, not just some machine churning out research. He cared about us as whole human beings, trying to understand our own hopes, joys and sorrows, exhibiting genuine concern for the things that were important to us both within the project and outside it. Whether we were feeling unmotivated to work or having trouble with qualifiers, or having a career crisis and deciding that you really wanted to study neuroscience, Allen patiently worked to accommodate us, always seeking first to further our own personal growth.

I would like to close on a personal note if I may. I will never forget the day I first met Allen Newell. It was on the 63rd floor of the U S Steel Tower, at the an-

nual departmental reception. The city lights glittered through the glass wall, members of the Pittsburgh Symphony played in the background, and as we talked, fireworks from some downtown event lit up the sky below us. I was naturally thrilled. But I had no reason to expect the thrill would last beyond the moment. I had no reason to expect what the next five years would bring. The impact Allen has had on my life, in some many different ways, is inestimable. I will be forever grateful for the time he spent with me and the confidence he showed in me.

Over the past year or two, I had the opportunity to give Allen a ride home on many occasions. I would usually pull around the back of his house while we continued talking for a minute. I will always remember him sitting in the car with his satchel, looking a bit cramped. My small Honda barely able to contain his energetic gestures. And that extra minute unfailingly turned into twenty or thirty until eventually I became worried that Noel would come out and get after us both for having to sit like that too long. Our conversations would range from technical discussions, to motivational sermons, to his concerns about the project, to historical anecdote that usually began — *is so and so a familiar name?....well this character.* Always in the end making a point for my benefit. And during these time and many others, I learned a lot about Allen. I learned how deeply he cared for the project, his students, his family. I learned how inseparable to him were the scientific process and the people involved in it. And I learned how much effort he put into both. How carefully he went about ensuring that we were reaching our potentials, that everyone had the opportunity to contribute and share in the excitement, that everyone felt they could come to him and honestly discuss their concerns.

I think only now I'm beginning to realize the significance of that effort. This is the thought I want to leave you with. It is perhaps easy for us to forget that Allen was human, too, that he went through periods of feeling less motivated, that he was anxious at times about how the world would receive his research, that he simply got tired. We see the hard work but forget that it came with ef-

fort — not just the science, but the relationships he built. Allen extended himself to us for our betterment and that is the essence of love. I'll miss him more than I can say.

Look at what has to be done [for Soar]

Have the right architecture

Learn continuously from experience

Communicate with the external world easily

Learn continuously from its environment

Live a long time

Embody all its tasks simultaneously

Become very large— 10^5 to 10^6 associations

Have a sense of history and place

Have a system with a sense of self

Learn from a social community

But Soar is not a final project—Simply a next project

A. Newell, 1991

Raj Reddy

I first met Allen in 1964. He was giving a talk at Stanford. The topic was GPS, the General Problem Solver. I said to myself, surely this is an impossible task. But as Allen began to explain the concepts of goals, operators and mean sense analysis, it was clear to me that he had captured the essence of the problem solving process. Daring to work on the hardest problems of science, that science has to offer, was the hallmark of Allen Newell.

I came to CMU in 1969. The main attraction was the exciting prospect of working with Allen, and I have never been disappointed. The environment, the values, the approach to quality of life we have today in the School of Computer Science are all directly inherited from Allen. Allen always gave freely of his time. By giving his complete attention to your problem, he made you feel very important. In each case, he would bring his unique talents of clarity of thinking and insightful understanding to the analysis of the problem at hand. Each time you would wonder how he could possibly have such deep insights on a problem that you had been working on for many months.

The greatness of Allen was he made everyone feel special. It's been our good fortune to have had Allen and Noel as our friends for many years. This whole week I went through the motions of work, there being this constant sense of loss. I cannot imagine our environment without Allen...

...project structures A through G were the only ones that suggested themselves to me as being plausible or interesting, but I was missing an obviously good bet in H, which was suggested by Allen Newell during a telephone conversation....

Structure H is based on a suggestion by Allen Newell that I consider to be an excellent one. It is that the "circle" kind of organization be supplemented by something that might be called a "cooperation expediter." Assume that there is a circle of contractors and, in addition, a single person — not so much a director as a facilitator or expediter of cooperation — who has control of enough money to put into effect on short notice various cooperative or supporting plans that are agreed upon by him and two or more members of the circle...

*Speech Understanding Systems: Final Report of a Study Group, A. Newell et. al., 1971,
section editor: J.C.R. Licklider*



Jill Fain Lehman

One way to describe the essential problem is that we are here touching one of the great scientific mysteries — the nature of intelligence. We are approaching it through a particular avenue, constructing intelligent machines, and sampling the whole at only a few places. It will be some time before enough of the picture accumulates to really see what is essential. There are whole stretches of intelligent action that have never been attended to in any scientific detail. Related to this is that the major source of knowledge about intelligent action is man himself and how he thinks.

Another way to describe the essential difficulty is that the nature of symbolic reasoning is such that it is easy for one mechanism to masquerade as another (the essence of simulation). It is possible to take an approach that has a grain of truth and force it to yield additional results, though with increasing difficulty. Alternative candidate mechanisms for intelligence can coexist for quite a while before resolution occurs, for each can be made to accomplish some of the functions of the other...

—A. Newell, W. Wulf, 1976

I am Jill Fain Lehman and I was a student and a colleague of Allen's. There are a set of moments spanning the last ten years that are, for me at least, quintessentially Allen. They all begin the same way. Allen and I are of course discussing work. I am sitting and most often he is standing back against the wall or leaning over the podium in his office. I just say something and the moment begins to unfold in the silence as I wait for his reply. Allen turns his head away to stare across the room or out the window, while he shakes his head once and says, *Now that's interesting*. What happens next changes across the moments, but there is something in the pause before he speaks, and in that phrase that tells me that I've caught him, and that whatever comes next he will push me down some path that I haven't thought about before...

I can't tell you the details of the different conversations that follow those moments — just that for the next few minutes or hours, there was nothing but the ideas and the excitement of the ideas and a chance to work out the ideas with him.

Large computer systems can be constructed by joining together many minicomputers — creating what can be called multi-mini-processors. The first such systems are just reaching the point where problems of programming and use dominate problems of design and construction...

We are not the architects of the multiprocessors we will describe. We are not even the primary systems programmers, who create the operating system and operating environment within which the user operates. We are users of the system. But we are not arms-length users, as are the users of a typical university computation center. For to use such a system one must indeed create a special programming system on it. Thus we are, shall we say, systems exploiters. We are just coming deeply into contact with our multiprocessors. We find ourselves facing many issues of how to exploit the system and to program it — of how to make it yield to our will...

One might say there is no issue — simply use the machine. But the question is not laid to rest so easily. Different strategies of exploitation require that effort be spent in different ways, thus precluding following alternative paths with any efficiency. Indeed, the issue as posed makes it sound like the multiprocessor arrived *sui generis* with the question of use fully open. That is not the case. The exploitation strategy is chosen before the design even begins and effects many of the structural features of the system. The actual situation is more like making a movie. Constructing the hardware system is like filming. Using it is like producing the movie in the editing room. The final editor is free to make any kind of movie he wants, but he must work with the film given him by the director....

...when completed with operating system and user facilities such as file systems and language processors, the system will look no different to the user than your local computation center. Only down in the boiler room, so to speak, will the multiprocessor design become apparent...

We might point out to psychologists that the problem is in essence faced by a population of intercommunicating humans. No one has internal symbols (i.e., addresses) designating all the things that all individuals designate internally. That is, they do without large addresses in the hardware. Instead they use language, which is a set of software-maintained large addresses, for their intercommunication. They continue to think their private thoughts in separate representational worlds. Thus the problem of communicating with small addresses is a fundamental one not restricted to the world of multi-mini-processors.

We have attempted to expose a set of programming issues that we have encountered in beginning to use a multi-mini-processor. We confess our fundamental ignorance of the correct formulations — to say nothing of the correct solutions — for most of these problems in this new environment. Perhaps these will no longer look like the important problems after we obtain more experience. That experience is now enveloping us day after day.

— *Some Issues in Programming Multi-Mini-Processors*,
A. Newell and G. Robertson, 1975

Bonnie E. John

I am Bonnie John and I was Allen's last non-Soar graduate student. When I met Allen in 1983, it was very easy to be intimidated. Not only because he was such a great intellect, but because he still had that office on the fourth floor with the big window behind his desk. And I'd come into the office to find this huge, booming silhouette against the half-open blinds. But it didn't take long for me to feel close enough that when I walked in the door, I'd flip on the lights and insist that he shut the blinds, so that I could have a discussion with a scientist and not the hulk.

About two and a half years into my Ph.D. program, I had to bring some work over to Allen at the house. And this was the first time I had been to the house, so I was very curious, of course, to see how this great scientist was at home. He opened the door and I almost fell backwards in shock, because here he was with a white t-shirt, painter's pants, some sort of blue and white kimono-type thing and a bandana tied around his neck. Very sharp contrast to the knit shirts and grey slacks that I could only categorize as 'old man clothes' that he wore to school. But he ushered me in and he told me to take a seat because he was on the phone with his son, and he'd be done in a minute. So I was listening to the one side of the conversation and he says, *I suppose you've heard about the honorary degree the University of Pennsylvania is going to give me and Yea, we'll go pick it up...* And then a little while later he said, *Yea, I'd gladly give up that degree if the Steelers had only won last week.* And it just hit me with such force that this great scientist was a whole person, with interests and priorities and personal feelings that had nothing to do with Soar or GOMS. But I

really should have known that, because I should have noticed much earlier that this whole person just shined through. And I can only describe it as beaming. You can see it in the picture on the program.

I remember him beaming at Herb when he read that just-so story about the djinn on the banks of the Monongahela River. And he beamed at me the first time I brought my baby girls into school — as well as when I got a consulting contract or had some sort of scientific breakthrough. And he beamed at each and every speaker at the Soar workshops as they told us about their contributions and progress in Soar. And this whole person just kept beaming through all the time.

Two years ago, when Allen told me that he had cancer, he said that he hoped that the news wouldn't slow things down, that we would all just keep on working. I reminded him that he had been living with this disease for almost a year already and that he had some time to come to terms with it and that we all had to come to terms with it. We had to have our time to grieve. And he stopped, and he beamed at me, and he said, *Of course you are right. We can do with expressing a little more love around here.*

And I think we did that in the last couple of years. And *of course* we're going to keep on working. And *of course* we're going to continue to advance his theories and our own. But Allen taught me not to forget the bigger picture. The whole person that allows us to do good science. And so, as he wanted, I'd just like to remind us all to keep expressing a little more love around here.

L^* is a system for building software systems. It is a tool for the professional programmer, and was originally intended for use in constructing artificial intelligence systems. Its most important use, however, has been in providing the basic software support for experimental computer systems. Under development at CMU since 1969, operational versions of L^* have existed since 1970 and have been in experimental use by a small community....

L^* 's roots lie in the series of IPLs, the original list processing languages (Newell and Shaw 1957; Newell, Tonge, Feigenbaum, Green, and Mealy 1964). As experience mounted with IPL-V and with LISP about the nature of system building in artificial intelligence, it seemed appropriate to make a fresh start with the emphasis on system implementation rather than on the language aspects. L6 (Knowlton 1966) had shown that efficient low level systems could be built using list structures as a data type. An early attempt to understand the lessons of L6 resulted in a similar macro-based system on the IBM 360 called *1 (Newell, Earley, and Haney 1967). An attempt to understand the nature of a flexible dynamic user interface resulted in a system call BIP (Basic Interface Package, Newell and Freeman 1968), embedded within IPL-V. All these systems can be taken as the direct precursors to L^* .

— *L^* : An Interactive, Symbolic Implementation System*,
A. Newell, D. McCracken, G. Robertson,
October 1977

...the role of minimally specified theories to keep design decisions open to accommodate new results in new task domains. This runs counter to some popular notions about rushing headlong for falsifiability. It also seems as if the theory were being adapted post hoc to meet each new result. But it seems a remarkably sensible strategy when a complete system such as the human brain has to be modeled. Far from being post hoc, there is an accumulation of design specification (i.e., of theoretical constraint) with each act of accommodation. As long as the theory has sufficient precision, all the chickens will come home to roost eventually. This tactic does lead to the temporary inability to predict behavior in various tasks, because of underspecification. PS.Harpy shows this in several ways. But this should just be seen as selecting an appropriate approximating sequence for developing a theory. After all, the empirical results that will cause difficulty are patient — they will lie in wait for a very long time.

—*Harpy, Production Systems and Human Cognition*,
A. Newell, 1978

John E. Laird

I am John Laird and I met Allen sixteen years ago, when I was deciding on where to go to graduate school. I had been going on a trip and I stopped in Pittsburgh, and as was custom, you stopped and talked to some of the faculty that you had heard of. I had never met Allen before but I had heard of him. So I scheduled an half hour meeting to discuss what was it like in Computer Science here at Carnegie Mellon.

So I went in and sat down and started to talk to him. And then he said, *Now, have you ever heard of production systems?* I said production systems? Sure, I know about producing things. I know about systems. I know about production systems. And then he sort of got out of me that I really didn't know what production systems were. And then he bounded to the blackboard and started to draw pictures of productions and working memories and things. And forty-five minutes later I was able to get out of there so I could get to my next meeting. And that's the way it was with Allen. He had this incredible enthusiasm about his work and was willing to share it with everyone. And he also taught you a lot of things at meetings like this. So I learned about production systems.

But I think the success in Allen's career is not from the things he produced or he told to us. It was in the questions he asked about our world and about us. It was his pursuit of the nature of mind and his ability to ask big questions. If you look at his research career, it is not one where he asked the question at the beginning, and then he followed up the ramifications, and then said, well how can I publish one more paper—get one more thing out of

this. It was he had this real innate curiosity about the nature of us and the nature of the mind. And whenever he came towards the closure of the limitations of one of the mechanisms he had developed, he then wanted to ask the questions about, *Now what's the next thing we can do with this?* or *How do we break this by looking at new phenomena?* So he was very good at producing the results, but the key to his greatness was in asking the questions. He would ask questions none of the rest of us would even think of.

As we have been working on Soar over the last ten years...it got to a point seven years ago I think where he felt things were going pretty well. It was time to ask new questions. And then he started saying, *Well, how do you emotions in Soar?* And this scared me. I said, "Well, wait a minute. Let's finish these things off." But he was willing to ask the new questions because he wasn't afraid of the truth. He wanted to find out what was the truth about us. What was the truth about the cognitive system? And so if we had to dash Soar along the way, that would be fine because he would learn some further truth.

But he didn't just do the asking the questions, he also did the details. And that's one of the things you see about his work. He always does a lot of the background work before he ever publishes his results. He told these great stories about when they were working on GPS. They had all these protocols of people solving cryptarithmetic problems. They got the graduate students and said, "Okay graduate students, here, go analyze these long protocols." And the graduate students

really were not too excited about that. So Herb and Al analyzed the protocols, and that's where their understanding came, and that's where their theories came. So it was the questions and then doing the details and then finally creating the theories. So the theories only came after that hard work.

And a final note that I'd like to say. Two and a half weeks ago, Paul and I came back to visit Al when he was in the hospital. We walked into his room and he was on his bed. It looked like he was asleep. And we came in and said hello. And then we started to talk about work. And suddenly his eyes brightened and we had a technical discussion for two hours about the future of Soar. It doesn't matter what the results are. Except for us, based on that decision, we made a dramatic left turn or right turn, as you want to think of it, in the directions of Soar.

But that was the way he was. He always had the enthusiasm for his work — willing to put things aside, even when he was sick, put the weaknesses and pain aside to contribute to the work because he wanted to get the answers. So that's where I'll stop. And that's where I'll always think about Al, is looking for those answers to these great questions.

Much remains to be specified in the above scheme to yield a definite information processing system. What happens (a likely occurrence) if more than one production is satisfied at once? What is the actual scheme for encoding information? What sort of collection of data structures constitutes the current state of knowledge on which the system works? What sort of tests are expressible in the conditions of productions? What sort of primitive operations are performable on the data and what collections of these are expressible in the actions of productions? What sorts of additional memories are available and how are they accessed and written into? How is the production system itself modified from within, or is this possible? How much time (or effort) is taken by the various components of the system and how do they combine to yield a total time for an entire processing?

There are many questions...

— *Production Systems: Models of Control Structure*
A. Newell, 1973

Paul Rosenbloom

Hello. I am Paul Rosenbloom. Over the last sixteen years Al has been my advisor, mentor, collaborator, friend and pretty much a second father. I'm afraid some of the things I'm going to say are going to duplicate what other people have said. I can see, in some ways, he strikes everyone in exactly the same way. But it's also been interesting to watch the evolution of how he strikes people over the years.

Discovering the mind was clearly Al's life, at least as we knew him. He went after it with a combination of intensity, enthusiasm, energy, commitment, confidence, selflessness, integrity, passion and joy, mastery and insight, which we may never see again. In my mother's terms he was a real mensch. It is an impossibility to convey a complete sense of Allen in only a few words. But let's see if I can illustrate some of this with a few anecdotes.

One that I will never forget was when he first told me that whenever he got tired of work he would just switch topics and start happily working on something else. I don't think it would have ever occurred to him to goof off. To him, science was relaxation. I also remember when some years ago he was dogged by two different ailments at the same time. One which made it hard for him to stand for long periods of time, and the other which made it difficult for him to sit for long periods of time. So what he did was to stack one terminal on top of the other and install a switch. So whenever he got tired of standing he could sit down and switch terminals. Whenever he got tired of sitting he could switch again and use the other terminal.

Probably the one that made the biggest impact on me was when John Laird was considering leaving

Xerox and hadn't yet found a position in which he was assured to be able to continue working on Soar. Al was aghast. How could anyone think of giving up research on something so important for any reason. In his words, *It would be a crime against humanity*. That's a direct quote. John did eventually find such a position at the University of Michigan. But the conversation comes back to us with a particular poignancy at a time like this. For Al to have been taken away from all of us — his family, those of us who worked closely with him and loved him, Carnegie and the School of Computer Science, the intellectual fields he helped establish, contributed so much to and had so much more to contribute to, and yes, humanity — it's a crime that has saddened all of us deeply.

It's hard to believe that I will never again get one of those handshakes that were complementary to the beaming. They were his way of showing when he thought we had really done a good piece of work. Or have a draft paper completely covered in red ink. Or be able to send him email at almost any time of day or night and almost immediately get back a very long and thoughtful response. I keep having this strong desire to cc: him on messages about Soar. And though I know it is crazy, I have this fantasy that all we need to do is to figure out how to lay a T1 line to the afterlife. I'm sure Al would be the first one to sign up.

It's not usual to draw on a person's faults at a time like this. But often their faults can be as illuminating, if not more so, than their strengths. With Al I only came up with three. His first fault was that he published too many ideas per paper and too many

papers in obscure places. The important thing was to get the ideas out, and seizing whatever opportunities arose to do so was what he did. Playing the academic game of tuning publications to maximize impact was never a significant concern. However, as a result, many of his best ideas got buried in the literature. There's a gold mine out there, waiting for someone with the initiative to go after it.

The second fault was that he had horrible taste in syntax. He was usually first with a new idea and first to put it into practice. But the resulting systems were often so ugly that they would then become superceded by a later system that embodied the same idea, but looked better — witness IPL and LISP, OPS and PROLOG. And I must say I'm keeping my fingers crossed about Soar.

The third fault was that he set such a high standard by his own behavior that you couldn't help feeling inadequate by comparison. But he inspired you to go beyond yourself in trying to achieve some part of that image and he always treated you like an equal. He even treated new graduate students like colleagues and thus taught them how to be so.

Al touched people deeply. Not necessarily through the ordinary means of being involved in their personal lives. I think he may have even, in fact, changed on this over the last few years. But for us, it was through his insights, his example of what a scientist should be like, and his loyalty to the people in institutions around him. His insights altered the courses of a number of scientific disciplines, and his people are telling me now — since they can't tell him — the research careers, both of people who knew him well and those who knew his work but barely knew him at all. His example inspired and instructed all those who came in contact with him. His loyalty was never more apparent than in the last year and a half, a significant portion of which went towards assuring that the Soar project would be able to continue as a vital enterprise and that the people who depended on him would be taken care of.

I know that ultimately the best way we can possibly honor Al, and the way he would have preferred, is not by these few inadequate words

spoken here, but by continuing to push the science towards a complete and integrated understanding of the mind.

For computer scientists, the immediate impulse for how to get something for nothing is to recurse—deciding the first step is just another task. Alas, the impulse is in vain. It is reminiscent of the old Pat-and-Mike story. Pat offered Mike a million dollars a year just to solve problems for him. That seemed to Mike a fantastic deal, and he accepted with alacrity. Knowing Pat to be a man of modest means, Mike said, "Where will you find all that money to pay me?" "That," said Pat, "is your first problem."

Formulating the Problem-Space Computational Model, A. Newell, G. Yost, J. Laird, P. Rosenbloom, E. Altmann, 1990

Stuart K. Card

I am Stuart Card. I came to CMU to study with Newell and Simon some years ago. It worked out better than I had ever expected. For the next twenty-five years I worked in a more or less, continuous association with Allen Newell. Science is partially learned through apprenticeship and mentoring and Allen performed that duty for me. I had an undergraduate background in Physics, similar to his, and he could initiate me into the strange worlds of psychological theory from that point of view. Allen Newell believed that science was not an occupation, but a calling. He had his problem, the nature of intelligence; and he had his method. He basically wanted an equivalent to differential equations that would allow the apparent complexity of behavior in the large to be derived from the playing out of the locally described mechanisms. IPL, LT, GPS, MERLIN, SOAR...they were all just refinements on the same basic scheme. Each one improving on the last.

Soar helped him to appreciate this subtle intertwining of problem solving and learning, to the point where he couldn't believe in a theory that didn't have both of those two intertwining. He had these incredible theories he was nursing along in the back of his mind, far beyond the current Soar. A theory of emotion. A theory of working memory. A theory of how regularities in the world induce a sort of furniture of the mind. For the last several years, every week on Thursday night, I used to talk to Allen from 11:00 to maybe 1:00 or 2:00 in the morning. I rearranged my kids' music lessons around this. Goodness knows what I made Noël rearrange on her side of the bargain.

My notebooks, like many of yours I suspect, are

filled with so many unfinished scientific discussions, so many things I still wish I could ask. In a way Allen Newell, this student of heuristic problem solving, lives on in his many heuristics and maxims for how to do science. Start by understanding the behavior, for example, or avoid ill-defined metrics like complexity.

That was Allen Newell the scientist. But there was also Allen Newell the organization man. Naturally, he also had maxims here, too, like never surprise your boss, or the purpose of a meeting is to ratify in public things previously negotiated in private. While managing to avoid the one thing he most feared in life, becoming a dean, he managed to have a strong hand in institution building. The Computer Science Department seemed to come into being. It grew satellites. It grew into a School. The computing infrastructure changed. Things happened to DARPA. Things happened to Xerox. Once at Xerox he helped me engineer a protective organizational structure around our joint project. At one critical meeting — at one critical juncture — he called a meeting with my boss, my boss' boss, my boss' boss' boss, and all of the department heads. To my amazement, by the time we walked out of there, we had this mandate to do whatever we wanted. And they all seemed so *happy* about it.

But it was from Allen Newell the human being and friend that I learned the most. He had this constructive optimism of life. You might describe it as universal subgoaling with chunking turned on. That is, he always preferred to focus his attention on what could be done about a problem than bemoan the fact that it existed. This was even true during his illness.

Perhaps the quality I most admired was the way he could combine warmth and integrity with force of action. He could cause an organization to do something without being Machiavellian or without using guile. He could be a tough negotiator, but straightforward and fair. How did he do that? It went beyond the skill to a way of being. I spent some twenty years admiring that property and trying to study how he did it.

But the thing I'm most grateful for was having Allen Newell as a great teacher and a great friend, for the many hours he generously gave me. Scientist, organization man, and human being. The best compliment I can give him is to use a word from the private Newell lexicon that we all know. Allen Newell was a man of *parts*.

There is no substitute for working hard—very hard.

—A. Newell

David Klahr

My name is David Klahr. Allen Newell influenced and enriched my life in many ways and for a long time. When I was still an undergraduate at M.I.T. in the late 50's, I came across a couple of odd papers by some folks named Newell and Simon. Several years later, while a graduate student at GSIA, I had the good fortune to attend a couple of Al's lectures on protocol analysis. Little did I know then that I would be even more fortunate and come to know him as a friend and colleague and co-researcher.

If there is a sing' term for Al it is *big*. He had big ideas. He had a big mind. He had a big heart. He sent big email messages. And he loved long phone calls. What else was big about him? You've heard some things already. His vision, his entrepreneurial style, his self-confidence, his energy, his enthusiasm, his personal and institutional loyalty, his unselfishness. And as you can see in the photograph, he was a big man with a big smile on his face most of the time.

I've been thinking a lot about Al's unique influence on cognitive psychology and on the field reaction to his contribution. I think I finally understand it. Al did something that made many psychologists uneasy. He took his experimental findings seriously. What do I mean by that? I mean that as a theorist — perhaps the greatest theorist of cognition in the world — Al looked at the body of accumulated knowledge and he believed the results. Not just his own, but the full gauntlet of published experimental results. He didn't carp about design flaws here, or the wrong statistical analysis there. He took the published record more seriously than

many card-carrying experimental psychologists. He believed that the results in the best journals should be accepted and that they demanded an integrated theoretical accounting. And that is what he devoted the last ten years of his life to. That's what Soar is all about.

A memorial service is supposed to provide a few memories and preparing for a service like this causes one to replay dozens or scores of hundreds, depending on how well you knew him, of interactions and conversations. It is hard to chose a single one. But I'll tell you about a conversation we had a couple of years ago. We were discussing some departmental business. We got to the topic of faculty retention. I asked Al for his advice on the kinds of inducements that I should be providing to keep our best people at CMU — both our younger faculty, who felt very mobile, and our well established stars, the eminent scientists who are always being courted by other universities. What did he think was appropriate to give them? Allen's reply was characteristically challenging and very revealing, although I didn't immediately understand it. *You got it all wrong*, he said. *You're asking the wrong question*. "What do you mean," I asked. He replied, *Being in an environment is a two-way affair*. I guess I looked puzzled. He said it again, *Being in an environment is a two-way affair*.

And then I got it. He meant that in order to feel attached or loyal or bonded to the place where you work, you have to see your impact on it. Allen knew that to really derive satisfaction from your university, you have to identify or create opportunities to contribute more than your eminence and

indirect costs. He found ways to do that, to work on the infrastructure, to build the very things that facilitated his teaching and research. He found ways to shape programs and departments and colleges and to nurture other peoples' careers. He was rightfully proud of many of the features of the CMU environment because his personal stamp is everywhere.

Being in an environment is a two-way affair. Allen Newell lived by that rule and I think it made him very happy.

ZOG is an interface for man-computer communication currently under exploration at CMU... ZOG may be described as a rapid response, large network, menu selection technique. The user sits in front of the video terminal on which is displayed a frame of information... Thus communication from the computer to the user is via a display of formatted alphanumeric data. Communication from the man to the computer is via the selection of options (some of which are called pads). Selection is accomplished either by touching the screen at the place where the option is displayed or typing the selection-character that appears at the front of each option or pad. Selection generally produces another frame of similar format with additional options. Selection may also evoke an action, which may lead to additional data being displayed or other activities taking place. Though it will not be of much concern here, ZOG is actually a communications device to other programs on the computer; thus the actions that it takes are to communicate commands and data to other programs (of arbitrary character), which in turn output back through ZOG to the user. From the user's viewpoint ZOG might as well be simply a program capable of carrying out complex activities...

What distinguishes ZOG from a standard menu-selection scheme, of which there are many, is that the response time for the next display is essentially instantaneous (eg, around a tenth to a quarter of a second) and that the total set of frames through which one can course is very large (eg, tens of thousands). These two features go together since rapid response in a small net is of only limited utility. With ZOG, the user is to be able to stay within the net essentially indefinitely, gaining knowledge and giving commands... Rapid response coupled with the large network produces a qualitatively different man-computer communication philosophy from standard menu selection... It is our intent to attend seriously to how to describe human performance in ZOG.

— *Notes for a Model of Human Performance in ZOG*
A. Newell, 1977

A. Nico Habermann

Thirty years ago I was a research assistant of Professor Adriaan de Groot in Amsterdam. We were doing a project in the study on the methodology of teaching mathematics. He went away in this four year. He went away for a whole year to the United States and one of the places that he visited and the people that he visited were Herb Simon and Allen Newell. This was the first time that I heard about them and about their work and about the idea of artificial intelligence. I was intrigued by it.

Eight years later I came to the United States as a visitor. And one of the first places I visited was on the invitation of the faculty here at Carnegie Mellon University to come to CMU. I was greatly impressed. I visited other places, went back to Holland, and came back permanently on the faculty in 1969. I never regretted it. I did not regret it because of the great minds and the great people here at Carnegie Mellon, in particular Allen Newell, Herb Simon, and others.

During those years, I have built a very close relationship and friendship with Allen. Many of you have already referred to his enthusiasm for science and the work that he has done. I was able to observe him while he was preparing his William James Lectures that he gave at Harvard, that ended up in the book that he recently published¹. I was really impressed with his enthusiasm, with his way of dealing with so many things at one time and being able to concentrate and publish such a great

work. I had other interactions with him and Herb Simon already mentioned his interest and concern about the Carnegie Mellon organization, in particular, the Department and later the School of Computer Science.

I was indeed one of those who had that privilege, to have him as my counselor on so many occasions and so many opportunities. He helped me think through the formation of the School of Computer Science in many ways. I also want to remind many of us of the great way in which he behaved at our Black Friday meetings and our faculty meetings. He was always a source of inspiration, one who was involved with the individual and concerned about the well-being of the individual students. He was always able to counter-argue our sense and need for regulations and trying to overcome our rigid judgments with the personal element.

I also had the privilege to see that he was honored for the work that he did. Not only was he honored here at CMU, by his students, by our faculty; he was also honored outside of our community. Herb Simon mentioned that he just received the National Medal of Science given by President Bush, that Paul received [for him] just about two weeks ago. I had the privilege of being there in Groeningen, the place where I was born in Holland, when he got an honorary degree in 1989. It was a great ceremony and it was also visible, but not only to CMU — for the world was really honoring Allen Newell.

¹Unified Theories of Cognition



Paul Allen Newell

I am his son, Paul. On behalf of my mother and myself, I wish to thank everyone who has been so helpful during Allen's illness and during the past week. Though we may have seemed overwhelmed, the caring and the offers of support are greatly appreciated. We also appreciate your presence here today — the CMU community, those who came from out-of-town on short notice, our family friends, my own friends.

My father ventured his own version of a scientific career in his last public talk entitled, *Desires and Diversions*, the one that Herbert referenced earlier. Mine was a different path in life. Problem spaces are just a variation of parallel parking in my world. We joked of the common threads — with both of us being on email, both of us following the Steelers, both of us knowing little about the other's profession. I had the upper hand, as commercials I had worked on would appear during the football games. Rarely did Chuck Knoll discuss cognitive psychology.

Underneath these jokes was the more important common thread of the respect both of us had for each other's endeavors. I guess in this community it was known as the handshake.

After Allen gave the *Desires and Diversions* lecture, we talked about it from many angles. The science involved to the actor's performance. We knew this might be an important talk from an historical perspective. We never thought it would be the last. So I have seen this tape many times and I wish to add a maxim that he and I defined during those sessions. To do great science you must have a supportive community, both professional and per-

sonal. Here at CMU, present and during all the time, all the years he was here, and with Noël, he had this.

One of his maxims from that speech was choose a final project to outlast you. He said he didn't think Soar would be that final project. But I think he underestimated the potential of Soar. Rather in that lecture, rather than talk from a global standpoint, I wish only to reflect on that maxim from a local view. Jill Lehman, Allen and I, along with other members of the Soar community, were working on a video of the Soar project. All the scientific quality of Soar with all the visual quality of Rhythm and Hues, a high-end, Hollywood computer animation house where I work. A merging of the best of our two professional worlds. After Jill and I had spent a couple of days pounding on the script and the storyboard, we met with Allen in the hospital in the beginning of July to review the latest tests I had done this spring and to plan where we wanted to be next. We will continue the project until the video is finished. His maxim is met.

On a more personal note, we spent many Sunday nights each fall solving the important problems of the world — the Steelers. Those of you who knew not to call him on Sunday afternoons were quite aware of the importance. He had played football in high school and he really cared about the Steelers and what happened to them. During the past couple of years he began watching hockey as well. With the Penguins as a local team, it's pretty easy to do that. I'm a baseball person. I tried to sell him on that sport. We went to a game last year, the first he had been to since he took me to Forbes Field in

the 1960's. During this last stay in the hospital, we watched as many Pirate games as possible. I think I was on my way to converting him. But above all, it was the Steelers that really mattered. And he fretted about them during the beginning of July and we were looking forward to many long distance phone calls this fall.

He had a maxim about football. This should not surprise anyone. I offer it to you with the request that it's a way of remembering Allen. You give the Steelers and their fortune your attention as well. He would pace back and forth in the upstairs study, nervous concern, until the maxim was met: *You can only relax when you have a three touchdown lead, it's the fourth quarter and you have the ball.*

I'm going to miss my father, not only as a father but as someone who I really enjoyed being around. I know that history will remember him as a great scientist. I hope it also records that his son thought he was a pretty great dad, too. Thank you.

The thoughts transcribed on pages 30-60 were offered by colleagues, friends and family at a Memorial Service in honor of Allen Newell, 27 July 1992, on the Carnegie Mellon University campus.

Fairy Tales

Allen Newell

Once upon a time, when it was still of some use to wish for what one wanted, ...

... there lived a King and Queen who had a daughter who was lovely to behold, but who never laughed.

Or perhaps:

... there lived an old fisherman by the side of a sea that had hardly any fishes in it.

If you are like me, you are already hooked. You are ready to abandon all talk of computers and electronic technology and professorships, and settle in to hear a fairy story. Their attraction reaches almost all of us.

They let us enter in upon an enchanted world. Magic abounds, though always in special ways. Animals talk, and not only animals but trees and bridges. Villainy is there, certainly danger. There are trials to be overcome -- usually three of them. But there is always the happy ending. The spell is broken and the Princess smiles and marries the youth who made her laugh. The old fisherman gets the Jinni back in the bottle with the top on. And happiness is ever after, which means at least for a little while.

The experts tell us that fairy stories are for childhood. That they contain lessons for the crises of growing up, and that their central attraction comes because they deal with what is central to this universal time of life.

Like Hansel and Gretel, we have to leave home and find our own way.

Like Jack, in the story of the beanstalk, we can bring home the bacon if we persevere, even if our parents don't think we can.

But there was more, if you remember your Jack. First he escaped back home with a bag of gold. But Jack and his mother used up the gold, showing that one success is not enough.

Then he made a second trip up the beanstalk to the Giant's castle. This time he came home with the magic hen that lays golden eggs, so he now had a technology for satisfying his and his mother's wants.

But even so, material things are not sufficient for the full life. So on his third trip Jack brought home the golden singing harp, symbolizing the higher things of life.

Or like the Princess with the Frog King, we must learn to keep our word and embrace what we find disgusting and ugly, to discover that it contains our heart's desire.

The experts notwithstanding, fairy stories are for all of us. Indeed, this is true if for no other reason than that today we are all of us children with respect to the future. We do not know what is coming. It is as new to us and as incomprehensible as adult life is to child. We find ourselves troubled and fearful at the changes taking place in ourselves and our society. We need the hidden guidance of fairy stories to tell us of the trials we must over-

come. To assure us that there will be a happy ending. Whether fairy stories have been written that speak to the heart of our own adult crises is unclear. How would we, the children, ever know? Perhaps we must get along with the fairy stories we have. We could do worse.

But even more, fairy stories seem to me to have a close connection to technology. That the aim of technology, when properly applied, is to build a land of Faerie.

Well, that should come as a shock! The intellectual garb of the modern academic is cynicism. Like a follower in a great herd, as surely as I am an academic, I am a cynic. Yet, I have just uttered a sentiment that is, if anything, straight from Pollyanna.

In point of fact, within the small circle of writers who manage to put technology and fairy stories between the same covers, the emphasis is always on the negative, on the dark side. The favorite stories are those that trouble.

Like the Sorcerer's Apprentice, who learns only enough magic to start the broom of technology hauling water from the River Rhine to the cistern, but who cannot stop it.

Like the Jinni in the bottle, where the story is never permitted to go to the conclusion in the Arabian Nights, with the Jinni snookered back into the bottle, but is always stopped with the Jinni hanging in air and the question along with it — Can we ever put the Jinni back? Or will there only be ink all over the sky 'til the stars go out?

Like the many stories of the three magic wishes, in which, promising infinite riches just for the asking, they are always spent, first on foolishness, second on disaster and third on bare recovery.

As in the Monkey's Paw, the old couple's first wish was for 200 pounds. That was foolish. The second wish was for the return of their just killed son. That was disaster. The third wish was to send their son back to his opened grave to try to recover for themselves a world where life could go on.

But I see it differently. I see the computer as the enchanted technology. Better, it is the technology of enchantment. I mean that quite literally, so I had best explain.

There are two essential ingredients in computer technology. First, it is the technology of how to apply knowledge to action to achieve goals. That is, it provides the capability for intelligent behavior. That is why we process data with computers — to get answers to solve our problems. That is what algorithms and programs are all about — frozen action to be thawed when needed.

The second ingredient is the miniaturization of the physical systems that have this ability for intelligent action. This is what Angel Jordan, my co-Whitaker Professor, has been telling us about. Computers are getting smaller, and cheaper, and faster, and more reliable, and less energy demanding. Everything is changing in the right direction together. The good things do not trade off against the bad ones. More speed does not mean more dollars. Smaller size does not mean lower reliability. On any given date, these tradeoffs that the economists so dearly love, of having to chose between painful options, certainly do hold. But come back next year and everything is better: smaller, cheaper, faster, more reliable, less energy.

Thus computer technology differs from all other technologies precisely in providing the capability for an enchanted world:

For brakes that know how to stop on wet pavements.

For instruments that can converse with their users.

For bridges that watch out for the safety of those who cross them.

For streetlights that care about those who stand under them — who know the way, so no one need get lost.

For little boxes that make out your income tax for you.

In short, computer technology offers the possibility of incorporating intelligent behavior in all the nooks and crannies of our world. With it we could build an enchanted land.

All very good. What about the Sorcerer's Apprentice? That comes about because of two half-fallacies. The first half-fallacy is that technologies are rigid and unthinking. Start the broom off carrying water and it does that and not something else. But every computer scientist recognizes in the Sorcerer's Apprentice simply a program with a bug in it, embedded in a first generation operation system with no built in panic button. Even with our computer systems today, poor things as they are, such blunderbuss looping is no longer a specter.

Exactly what the computer provides is the ability to not be rigid and unthinking, but rather to behave conditionally. That is what it means to apply knowledge to action: it means to let the action taken reflect knowledge of the situation, to be sometimes this way, sometimes that, as appropriate. With small amounts of computer technology — that is, with small amounts of memory and small amounts of processing per decision — you often can't be conditional enough. That is certainly the story of the first decades of the computer revolution. It was too expensive and involved too much complexity to create systems with enough conditionality. We didn't know how and couldn't have afforded it if we had. Consequently, many applications were rigid and unthinking. It was indeed a Sorcerer's Apprentice who seemed to run the computerized billing service.

However, the import of miniaturization is that ultimately, we will be able to have enough capability for conditionality in a small enough space. And the import of our scientific study of computers is that we will know how to make all the conditionality work for us. Then the brooms of the world themselves can know enough to stop when things go wrong.

The second half-fallacy behind the Sorcerer's Apprentice is that technologies by their nature extract too high a price. That is a message of the recent literature of political ecology. Our technologies inevitably demand that we use up our precious

world. There is rather abundant evidence for this view. Here in Western Pennsylvania, the price to the enchantment of our countryside from taking our coal by strip mining is only too much in evidence. Less in our awareness, because it was so thorough, was what the loggers did to Western Pennsylvania. Not once, but thrice, within forty years they swept the hillsides almost bare. The hot scalding breath of a dragon could hardly have done better for desolation.

But all is not inevitable. Ecologically, computer technology itself is nearly magic. The better it gets, the less of our environment it consumes. It is clean, unobtrusive, consumes little energy and little material. And as we push it to higher peaks of speed and memory, it becomes more of all these things. For deep technical reasons this has to be. There is no way to obtain immense amounts of processing power by freezing technology at some cost in dollars, material and energy per unit of computation, and then just buying more and more of it, consuming our wealth and our environment. Instead, for a long time to come, as we get more and more of it, the less will it impact our environment.

Even more, the computer is exactly the technology to permit us to cope intelligently with the use of our other resources. Again, by providing us with distributed intelligence, it can let us keep track of the use and abuse of our environment. And not only of the destruction that we ourselves visit on our world, but also that which nature does as well. Mt. Vesuvius was hardly bound by any antipollution ordinances posted on the walls of ancient Pompeii.

In sum, technology can be controlled, especially if it is saturated with intelligence to watch over how it goes, to keep accounts, to prevent errors, and to provide wisdom to each decision. And these guardians of our world, these magic informational dwarfs, need not extract too high a price.

But I said that the Sorcerer's Apprentice was guided by *half-fallacies*. I did not dismiss the view totally. Because, of course, in fairy stories there are great trials to be performed before the happy ending. Great dangers must be encountered and

overcome. Because also, in fairy stories, the hero — the one who achieves finally the happy ending (and it is as often a girl-child as a boy-child) must grow in virtue and in mature understanding. No villains need apply for the central role. The fairy story that I am indirectly spinning here will not come automatically and we must earn it.

Where are we now? We are not at the end of the story, though we are surely at the end of my talk. In fact the fairy story is hardly past its "Once upon a time". Still, I wish to assert that computer science and technology are the stuff out of which the future fairy land can be built. My faith is that the trials can be endured successfully, even by us children who fear that we are not so wise as we need to be. I might remind you, by the way, that the hero never has to make it all on his own. Prometheus is not the central character of any fairy story, but of a tragic myth. In fairy stories, magic friends sustain our hero and help him overcome the giants and the witches that beset him.

Finally, I wish to express my feeling of childlike wonder that my time to be awake on this earth has placed me in the middle of this particular fairy story.

Presented 17 September 1976 at an Inaugural Party for the establishment of the U.A. and Helen Whitaker Professorships at Carnegie Mellon University. This article originally appeared in FOCUS, December 1976.

Discovery-Integration Processes for Breaks

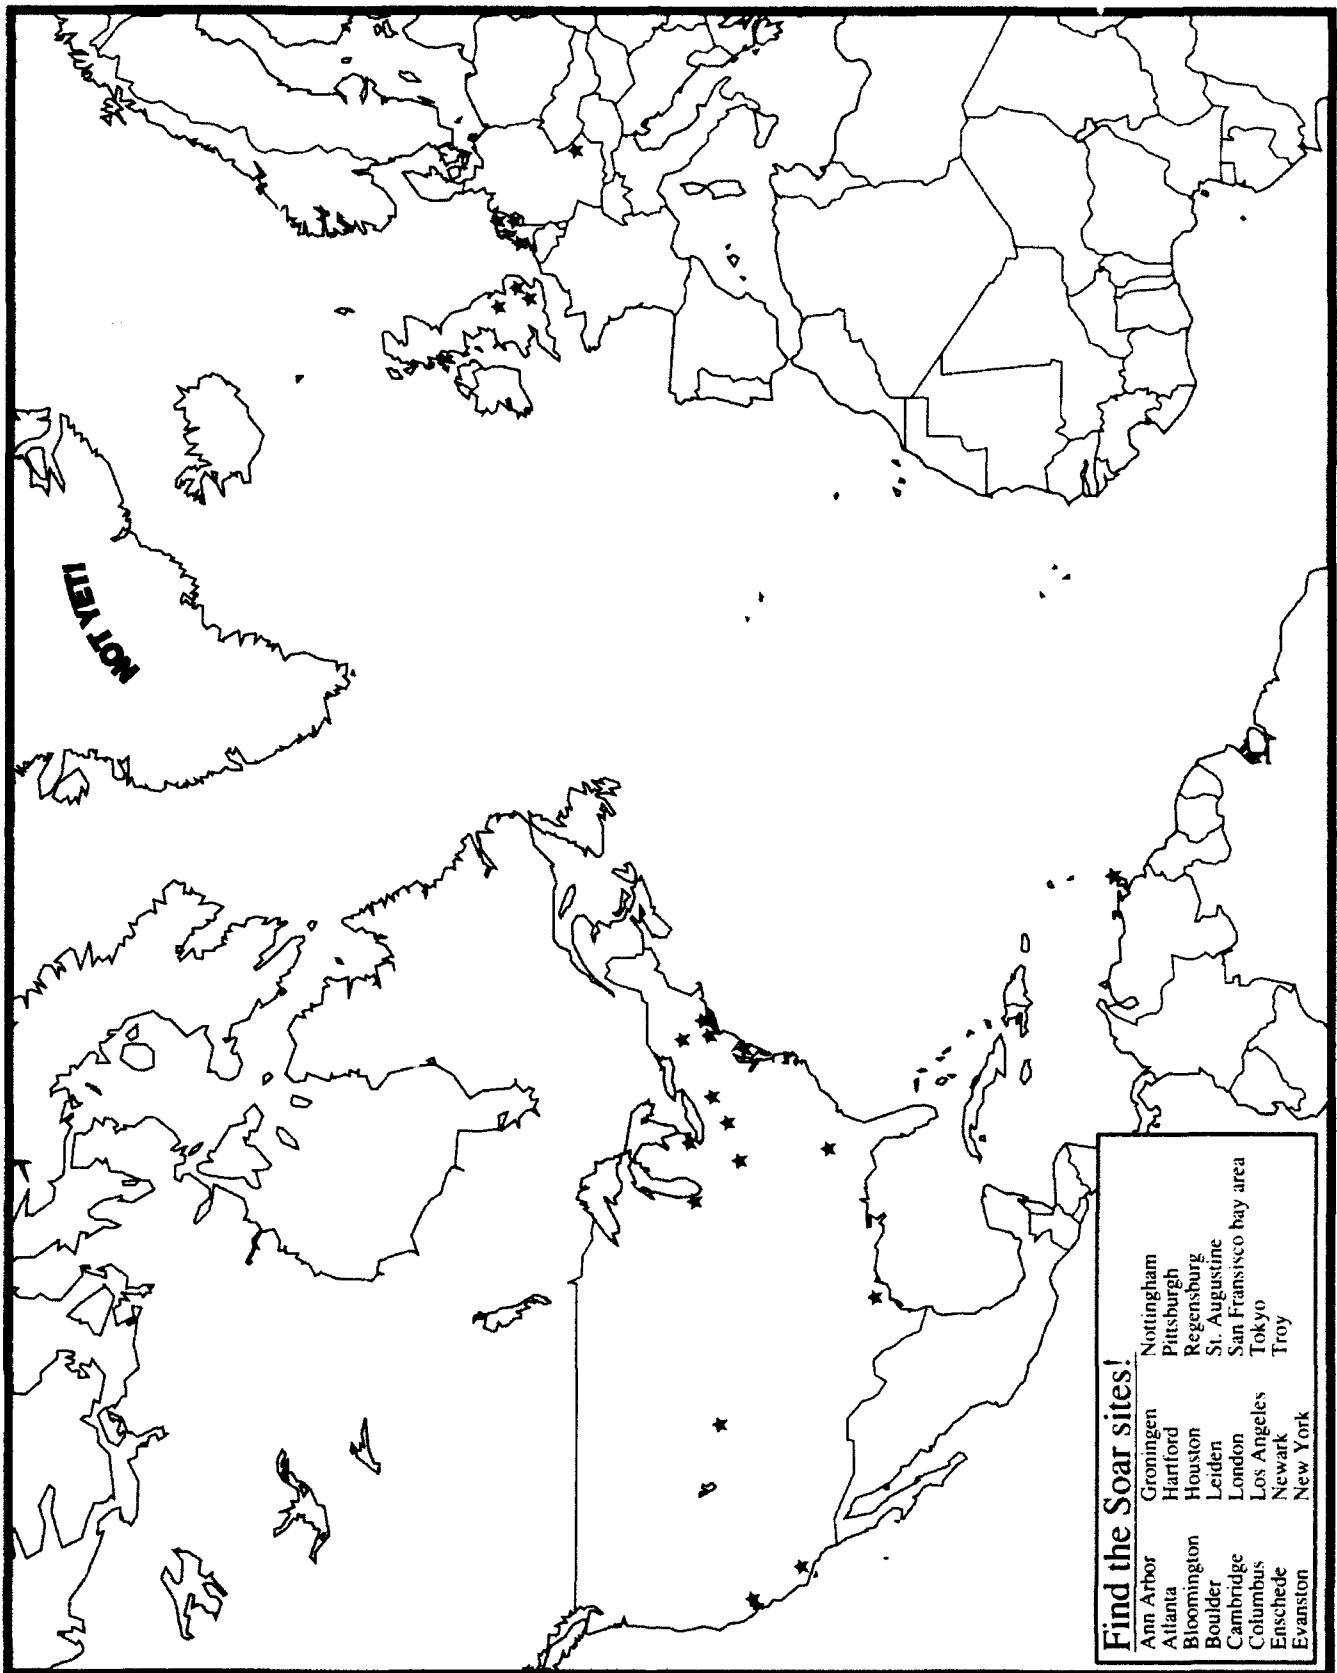
"The whole area of problem solving makes a pretty illustration of the discovery-integration process. Many separate problem solvers have been written; the tasks have usually been idiosyncratic and non-utilitarian — games, puzzles, and areas of elementary mathematics. Some of the individual results have been impressive. Out of these separate efforts has emerged a clear understanding of two related, important problem solving mechanisms: the formulation of a problem as the search through an exponentially growing tree of possible solutions; and the role of various heuristics to narrow and guide the search so it can be successful within reasonable processing times."

— A. Newell, A. Perlis, E. Schatz, 1964

Eaten all the doughnuts you can eat? Read all the biographies a dozen times or more? Chatted sufficiently to last the day? Finding your mind stimulated from the subject matter?

Then take these *break* moments to explore basic problem solving. We expect you will soar through these tasks. You will need:

1. *One Mind Matters* Pencil
2. *One Mind Matters* Work/Note Pad
3. A little time

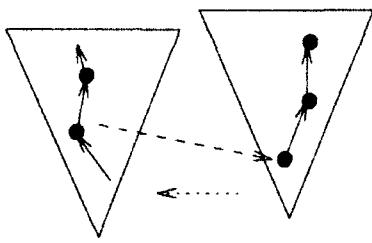


Find the Soar sites!

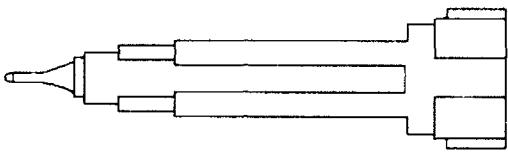
Ann Arbor	Groningen	Nottingham
Atlanta	Hanford	Pittsburgh
Bloomington	Houston	Regensburg
Boulder	Leiden	St. Augustine
Cambridge	London	San Francisco bay area
Columbus	Los Angeles	Tokyo
Enschede	Newark	Troy
Evanson	New York	

Can you name these...?

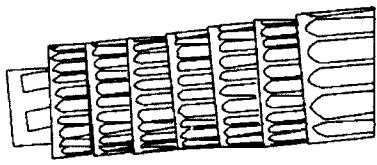
G R E A T M O M E N T S I N A R C H I T E C T U R E



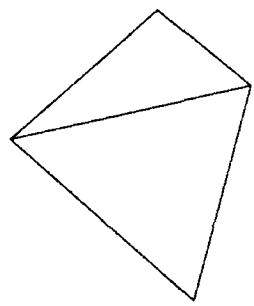
1982



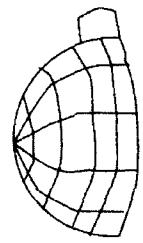
1931



1350



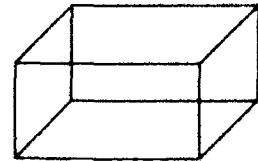
2500 B.C.



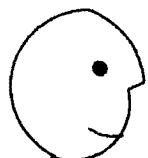
? B.C.

The following problem was presented to five cognitive scientists.
Can you draw a line from each response to the scientist who generated it?

Here is a four ounce block of jello. If you eat exactly one-half of it, how much will be left?

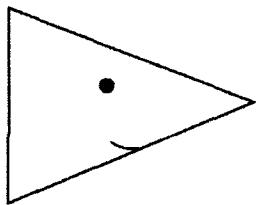


Situated
cognitivist



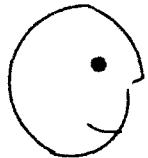
That reminds me of a story

Soar
researcher



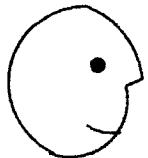
Does anyone have a knife?

Philosopher



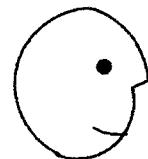
I seem to be facing the wrong way.

Case-based
reasoner



□ ■ □ ■ □
4 2 pounds ounces -ed

Connectionist



This actually can't be done. First, you'd have to slice off a quarter, then a quarter of what's left, then....

Find the Minds

Can you find the symposium speakers in this chart?
The names can be found horizontally, vertically or diagonally.

T O L K B I O L K L I M I N D M A T E R S R
O A R J O H N E L I M O O L B N E S O R T E
P S T U N O A W K S R E T T A M D N I M E T
E J O H N L O I N L M J G S T E V T I E I N
B O M B I K S T U R A J U R T H K L U H E E
M H R C E H A M I T R H D Y O U L O P R R P
A N T J A M R A R L A I R D K E A T H A D R
T S O W L R E R O B E R T A R B G R O B I A
T O J I L L D C D J L E Z V G O R D T I A C
U N N J E O D E D U S P E I R I A N N A L A
J L Z Y N P Y L Y S H Y N D Z E O N A N M I
F A I N T E H M I T E L O B J E S I M O N C
M I L I N D W W N E R N N U U R N M A T U I
E R A U L M E E Y I B O O N S O B O E J O R
L D P A U L I R L E E Y M I T C H E L L Y T
Y O U N G Z A S L L R R I C H A R D T S P A
R O S E N M L L J S T E S H A D M A T T E P

Gordon	Bell	Milind	Tambe
Stuart	Card	Bonnie	John
Philip	JohnsonLaird	Thad	Polk
Marcel	Just	Richard	Young
David	Klahr	Tony	Simon
Robert	Mehrabian	John	Laird
George	Miller	Jill	Fain Lehman
Tom	Mitchell	Paul	Rosenbloom
Zenon	Pylyshyn	Richard	Lewis
Mary	Shaw	David	Steier
Raj	Reddy	Allen	Newell
Herbert	Simon	Mind	Matters
Patricia	Carpenter	Soar	GOOD JOB

Mindful Matters

How many words can you find in the word: **SOAR**
[Soar found twelve]

How many words can you find in the words: **COGNITIVE SKILLS**
[The organizing committee found two]

.R AIQUIZ

"A program which generates a sequence of one minute questions proposed by Allen Newell and written by George Robertson is available on the CMU PDP-18 and can be activated by the following procedures by a teletype: .R AIQUIZ." These questions were used for testing the general awareness and breadth of a student in the AI area. This is only a sampling...

- Q: What is the current state of work on recognition?
- Q: Name a task that facial feature recognition systems does not do, but might do with some modification.
- Q: What have been the contributions of AI to physics?
- Q: Has there been any contribution of operating systems to AI?
- Q: What was the first AI effort on sequential concept formation?
- Q: What are some of the contributions that Green made to AI?
- Q: Has causality been accomplished by an AI program?
- Q: Is branch-and-bound a heuristic?
- Q: What is the current state of work on the concept of time?
- Q: In what task can operator selection not be used?
- Q: Could Stanford's hand-eye system be modified to do seriation?
- Q: Name some specific varieties of induction accomplished by hill climbers.
- Q: Has sequential concept formation been accomplished by an AI program?
- Q: Name some specific varieties of games accomplished by heuristic search programs.
- Q: What method does SAM use?
- Q: Has there been any contribution of AI to political science?
- Q: What gives GPS its power?
- Q: Compare work on relational concept formation and recognition.
- Q: Could SKETCH PAD be modified to do assembly line balancing?
- Q: Name a task that SAINT does not do, but might do with some modification.
- Q: What are programs that perform several tasks?

- Q: Who developed horizon phenomena?
- Q: Name another program other than Samuel's checker program that works on the same tasks.
- Q: What was the first AI effort on induction?
- Q: Where is AI research on induction going on?
- Q: What have been the contributions of physics to AI?
- Q: Has there been any contribution of chemistry to AI?
- Q: Is branch-and-bound a heuristic?
- Q: What are some of the contributions that Feldman made to AI?
- Q: Why is a task environment so called?
- Q: What is the current achievement in the concept of time?
- Q: Is depth-first search a heuristic?
- Q: What is resolution?
- Q: Why is a problem so called?
- Q: What method does SIR use?
- Q: What tasks does REF-ARF work on?
- Q: What is the main contribution of MATHLAB to AI?
- Q: What method does GPS use?
- Q: Name another program other than GPS that works on the same tasks.
- Q: Why is Weisenbaum's Nim player so called?
- Q: Is feature extraction used in BASEBALL?
- Q: What is the main contribution of GPS to AI?
- Q: What are the differences between GPS and STRIPS?
- Q: What are the main results of Minsky and Pappert's "Perceptions" book?
- Q: What are the current bottlenecks in building integrated robots?
- Q: What is the role of higher order logics in theorem proving?
- Q: Are there programs that modify their representation in solving a problem?
- Q: What is an intelligent action?
- Q: Name an intellectual task that no program has done.

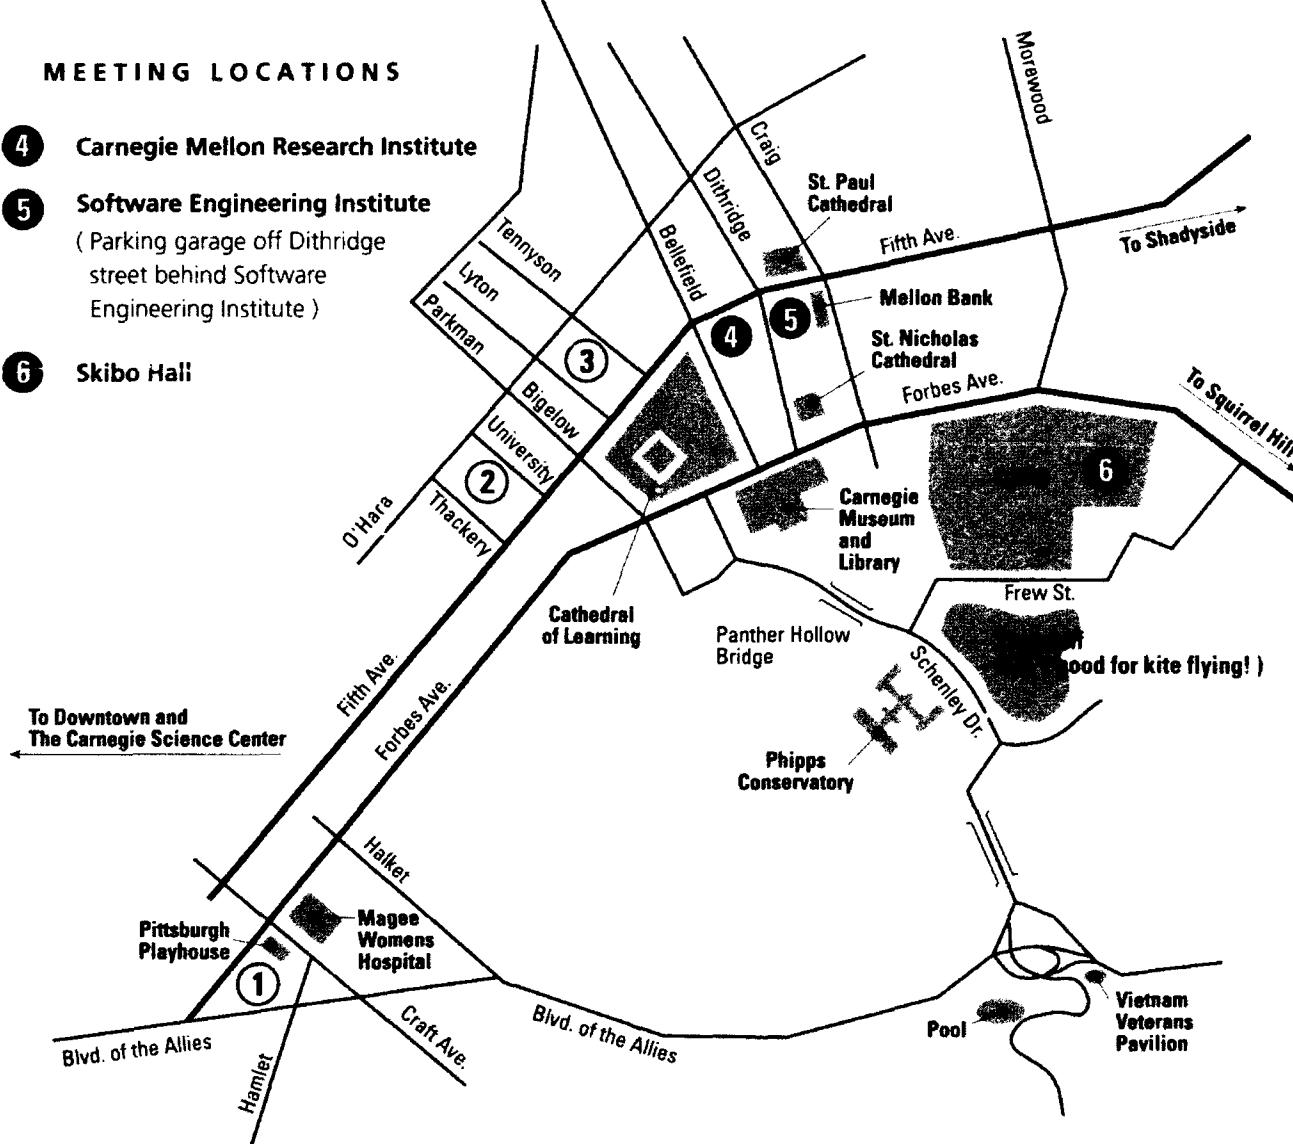
M A P

H O T E L S

- 1 Hampton Inn
- 2 University Club
- 3 Holiday Inn

M E E T I N G L O C A T I O N S

- 4 Carnegie Mellon Research Institute
- 5 Software Engineering Institute
(Parking garage off Dithridge
street behind Software
Engineering Institute)
- 6 Skibo Hall







With Gratitude

The School of Computer Science wishes to thank the following individuals and organizations for contributing their immense talents and generous support in preparation for Mind Matters :

Noël Newell and Paul Allen Newell

The National Science Foundation

The Office of Naval Research

Hugo T. Cheng and MAYA Design Group

Kenneth Andreyo: for his many fine photographs

Richard Lewis

James Kocher and Mary Jo Dowling, The Graphics Deli, School of Computer Science

Gabrielle Michalek, Carnegie Mellon University Archives

Our Distinguished Speakers

The Program Committee: Catherine Copetas, John Laird, Jill Fain Lehman, John McDermott, Tom Mitchell, Paul Rosenbloom, David Steier

Barbara Andreyo

Betsy Herk

Frank Ritter

The Soar Group

Instructional Technology, Carnegie Mellon University

The many, many friends and colleagues of Allen Newell

Soar Foods, Inc.: *"Exhausted? Nothing refreshes like a hot cup of Quiescence Tea."*

The proceedings from the Mind Matters Symposium will appear as part of
the *Carnegie A "on Symposia on Cognition* series, published by
Lawrence Erlbaum Associates, Inc.

Cover Design: Hugo T. Cheng, MAYA Design Group

Printing: Herrmann Printing and Litho, Inc.

Copyright © 1992 School of Computer Science

Contributions to Cognitive and Computer Science in Honor of Allen Newell